

# Object Storage Service

## User Guide

**Issue** 11  
**Date** 2023-11-24



**Copyright © Huawei 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 Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks 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 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.

# Security Declaration

## Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

---

# Contents

---

<b>1 Before You Start.....</b>	<b>1</b>
1.1 Overview.....	1
1.2 Using OBS.....	2
1.3 Accessing OBS.....	2
<b>2 Storage Classes.....</b>	<b>8</b>
<b>3 Bucket Management.....</b>	<b>14</b>
3.1 Bucket Overview.....	14
3.2 Creating a Bucket.....	15
3.3 Listing Buckets.....	17
3.4 Viewing Bucket Information.....	17
3.5 Managing Bucket Quotas.....	17
3.6 Deleting Buckets.....	18
<b>4 Object Management.....</b>	<b>20</b>
4.1 Object Overview.....	20
4.2 Uploading Objects.....	23
4.2.1 Uploading an Object.....	24
4.2.2 Uploading Objects Using a Multipart Upload.....	26
4.3 Downloading Objects.....	31
4.4 Managing Folders.....	32
4.4.1 Creating a Folder.....	32
4.5 Other Object Operations.....	33
4.5.1 Listing Objects.....	33
4.5.2 Copying Objects.....	34
4.5.3 Viewing Object Information.....	35
4.5.4 Sharing Objects.....	35
4.5.5 Managing Object Metadata.....	37
4.5.6 Restoring Objects from the Archive Storage.....	48
4.5.7 Direct Reading.....	49
4.6 Deleting Objects.....	50
4.6.1 Deleting an Object.....	50
4.6.2 Undeleting an Object.....	51
4.6.3 Managing Fragments.....	52

<b>5 Permissions Control</b> .....	<b>54</b>
<b>6 Data Management</b> .....	<b>55</b>
6.1 Lifecycle Management.....	55
6.2 Tags.....	65
6.3 Bucket Inventory.....	66
6.4 Event Notifications.....	71
<b>7 Data Access</b> .....	<b>79</b>
7.1 Static Website Hosting.....	79
7.2 Domain Name Management.....	87
7.2.1 User-Defined Domain Name Configuration.....	87
<b>8 Data Security</b> .....	<b>89</b>
8.1 Server-Side Encryption.....	89
8.2 CORS.....	90
8.3 Versioning.....	92
8.4 URL Validation.....	95
<b>9 Monitoring and Logging</b> .....	<b>98</b>
9.1 Monitoring.....	98
9.2 Auditing.....	99
9.3 Logging.....	100
<b>10 Parallel File System</b> .....	<b>104</b>
<b>11 Change History</b> .....	<b>105</b>

# 1 Before You Start

## 1.1 Overview

Object Storage Service (OBS) provides a series of documents to help you better understand and use the service.

Before using OBS, read this document and other related help documents (see [Table 1-1](#)) to understand the basic concepts, application scenarios, and operations of OBS, so you can quickly get started with OBS.

**Table 1-1** OBS documents

Document	Description
<a href="#">Service Overview</a>	Introduces the basic concepts, advantages, application scenarios, billing, and permissions management of OBS.
<a href="#">Getting Started</a>	Helps you quickly get started with OBS, such as creating buckets, uploading objects, and downloading objects.
<a href="#">Console Operation Guide</a>	Guides you through OBS console operations with elaborated use examples.
<a href="#">Best Practices</a>	Provides a collection of OBS best practices for typical scenarios, helping you achieve your business goals with better performance, lower costs, and more convenient operations.
<a href="#">API Reference</a>	Lists the REST APIs provided by OBS, including request samples, response samples, and parameter descriptions. It helps you use APIs to conduct secondary development.
<a href="#">FAQs</a>	Offers you the answers to frequently asked questions about OBS.

## 1.2 Using OBS

You can manage OBS resources in the ways listed in the table below.

Tool	Description	How to Use	Reference
OBS Console	OBS Console is a web-based GUI where you can manage your OBS resources with ease.	Create an account or an IAM user to log in to OBS Console.  For details, see <a href="#">Logging In to OBS Console</a> .	<a href="#">Console Operation Guide</a>
API	OBS provides REST APIs that support HTTP/HTTPS requests. You can call these APIs to create, modify, and delete buckets, as well as to upload, download, and delete objects.	Configure the endpoint during environment preparation and make the access keys (AK and SK) part of your signature in any requests you send. If you use REST APIs to develop a program, you need to calculate the signature based on the signature algorithm defined by OBS and add the signature to the requests.  For details, see <a href="#">User Signature Authentication</a> .	<a href="#">API Reference</a>

## 1.3 Accessing OBS

### OBS Domain Name

The following two concepts are related to OBS domain names:

1. **Endpoint:** OBS provides an endpoint for each region. An endpoint is a domain name to access OBS in a certain region and is used to receive access requests sent from that region.
2. **Bucket domain name:** Each bucket in OBS has a domain name. A domain name is the Internet address of a bucket and can be used to access the bucket over the Internet. It is typically used in cloud application development and data sharing scenarios.

An OBS bucket domain name is in the *BucketName.Endpoint* format.

*BucketName* indicates the name of a bucket, and *Endpoint* indicates the OBS domain name of the region where the bucket is located.

**Table 1-2** lists the bucket domain name, and other domain names of OBS, including their formats and protocols.

**Table 1-2** OBS domain names

Type	Structure	Description	Protocol Type
Region domain name	[Structure] <i>Endpoint</i> [Example] obs.eu-west-101.myhuaweicloud.eu	Each region has an OBS endpoint, which is the OBS service domain name of the region.	HTTPS HTTP
Bucket domain name	[Structure] <i>BucketName.Endpoint</i> [Example] bucketname.obs.eu-west-101.myhuaweicloud.eu	After a bucket is created, you can use the domain name to access the bucket. You can assemble the domain name by putting the bucket name and endpoint together, or you can obtain it by <a href="#">viewing the basic bucket information</a> on OBS Console or OBS Browser+.	HTTPS HTTP
Object domain name	[Structure] <i>BucketName.Endpoint/ObjectName</i> [Example] bucketname.obs.eu-west-101.myhuaweicloud.eu/object.txt	After an object is uploaded to a bucket, you can use the object domain name to access the object. You can assemble the object domain name by putting the bucket name, OBS service endpoint, and object name together, or you can obtain it by <a href="#">Viewing Object Information</a> on OBS Console or OBS Browser+. Alternatively, you can call the GetObjectUrl API through an SDK to obtain the object domain name.	HTTPS HTTP
<b>Static website domain name</b>	[Structure] <i>BucketName.obs-website.Endpoint</i> [Example] bucketname.obs-website.eu-west-101.myhuaweicloud.eu	A static website domain name is a bucket domain name when the bucket is configured to host a static website.	HTTPS HTTP

## Endpoints

OBS has an endpoint in each region.

Generally, the endpoint carried in a request for accessing OBS must be the endpoint of the region where the requested resource resides. However, in some special cases, you can use any endpoint.



## 1. Scenarios where the endpoint in a request must be the endpoint of the region where the requested resources reside

When accessing a bucket or an object, the endpoint in the request must be the endpoint of the region where the bucket or object resides.

For example, if bucket **mybucket** is in region **eu-west-101**, you can list objects in the bucket by sending a request shown in the following example:

### A correct example of request and response for listing objects:

#### [Request]

```
GET / HTTP/1.1
Host: mybucket.obs.eu-west-101.myhuaweicloud.eu
Accept: */*
Date: Thu, 10 Mar 2016 08:51:25 GMT
Authorization: authorization
```

#### [Response]

```
HTTP/1.1 200 OK
x-obs-request-id: 0001EF710C000001536176DA465E4E6G
x-obs-id-2: Rdj0zVzRkihRcjcQUqjkDGt8JuAgi2CGuLiP7Pv/cYYplsS0xTFJQHP5vSg5yOYC
Content-Type: application/xml
Date: Thu, 10 Mar 2016 16:58:12 GMT
x-obs-bucket-location: eu-west-101
Content-Length: 259
```

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<ListBucketResult xmlns="http://obs.myhuaweicloud.eu/doc/2015-06-30/">
  <Name>mybucket</Name>
  <Prefix/>
  <Marker/>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>>false</IsTruncated>
  <Contents>
    <Key>object001</Key>
    <LastModified>2015-07-01T00:32:16.482Z</LastModified>
    <ETag>"2fa3bcaaec668adc5da177e67a122d7c"</ETag>
    <Size>12041</Size>
    <Owner>
      <ID>b4bf1b36d9ca43d984fbc9491b6fce9</ID>
    </Owner>
    <StorageClass>STANDARD</StorageClass>
  </Contents>
</ListBucketResult>
```

If the endpoint in the request is not consistent with the endpoint of the requested region, an error message is returned indicating that the bucket does not exist.

In the example above, bucket **mybucket** is in region **eu-west-101**. If you use the endpoint of **cn-south-1** (**mybucket.obs.cn-south-1.myhuaweicloud.eu**) to access the bucket, HTTP 404 is returned, indicating that the bucket does not exist. In this case, you can call the API for [obtaining bucket location](#) to obtain the bucket's region ID, and then send the request again.

### An incorrect example of request and response for listing objects:

#### [Request]

```
GET / HTTP/1.1
Host: mybucket.obs.cn-south-1.myhuaweicloud.eu
Accept: */*
Date: Thu, 10 Mar 2016 08:51:25 GMT
Authorization: authorization
```

#### [Response]

```
HTTP/1.1 404 NoSuchBucket
x-obs-request-id: 0001EF710C000001536176DA465E4E6H
```

```
x-obs-id-2: Rdj0zZvRkihRcjcQUqjkDGt8JuAgi2CGuLiP7Pv/cYYplsS0xTFJQHP5vSg5yOYL
Date: Thu, 10 Mar 2016 08:51:30 GMT
Content-Length: 0
```

## 2. Scenarios where any endpoint can be used in a request

You can use the endpoint of any region in the API requests for listing buckets and for obtaining a bucket's region information, because these two APIs search for buckets from all regions.

### An example of request and response for obtaining a bucket's region information:

For example, if bucket **mybucket** is in region **eu-west-101** and the endpoint of the **cn-south-1** region is used in the request, the bucket location information can still be obtained.

#### [Request]

```
GET /?location HTTP/1.1
Host: mybucket.obs.cn-south-1.myhuaweicloud.eu
Accept: */*
Date: Thu, 10 Mar 2016 08:51:25 GMT
Authorization: authorization
```

#### [Response]

```
HTTP/1.1 200 OK
x-obs-request-id: 0001EF710C000001536176DA465E4E6G
x-obs-id-2: Rdj0zZvRkihRcjcQUqjkDGt8JuAgi2CGuLiP7Pv/cYYplsS0xTFJQHP5vSg5yOYC
Content-Type: application/xml
Date: Thu, 10 Mar 2016 16:58:12 GMT
Content-Length: length

<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<Location xmlns="http://obs.myhuaweicloud.eu/doc/2015-06-30/">eu-west-101</Location>
```

## Accessing OBS over the Internet

Accessing OBS over a public network may generate inbound traffic for write operations (for example, uploading data to OBS), as well as outbound traffic for read operations (for example, downloading data from OBS). Inbound traffic does not incur fees, but outbound traffic does.

If you access OBS over a public network, you can use a URL to specify resources in OBS. An OBS URL is structured as follows:

*Protocol://[BucketName.]Endpoint[:Port][Object][?Param]*

**Table 1-3** Parameters in an OBS URL

Parameter	Description	Mandatory or Optional
Protocol	The protocol used for sending requests, which can be either HTTP or HTTPS. HTTPS is a protocol that ensures secure access to resources. OBS supports both HTTP and HTTPS.	Mandatory
BucketName	Name of the requested bucket, which uniquely identifies a bucket in OBS.	Optional

Parameter	Description	Mandatory or Optional
Endpoint	Domain name (endpoint) of the region where the OBS bucket is located.	Mandatory
Port	The port enabled for protocols used for sending requests. The value varies depending on software server deployment. If no port number is specified, the protocol uses the default value. Each transmission protocol has its default port number. In OBS, the default HTTP port number is <b>80</b> and that of HTTPS is <b>443</b> .	Optional
Object	Path of the requested object resource.	Optional
Param	Specific resource contained by a bucket or object. The default value of this parameter indicates that the bucket or object itself is obtained.	Optional

Example: You have a bucket named **mybucket** in the EU-Dublin (**eu-west-101**) region. The bucket contains an object named **myfolder/myfile.txt**. The URL for accessing the object over the public network is as follows:

**<https://mybucket.obs.eu-west-101.myhuaweicloud.eu/myfolder/myfile.txt>**

#### NOTE

All API requests except the one for listing objects must contain the **BucketName**. In consideration of the DNS resolution performance and reliability, OBS requires that the bucket name must precede the **Endpoint** when a request carrying a bucket name is constructed to form a three-level domain name, also mentioned as virtual hosting access domain name.

## Checking OBS Version (OBS 2.0 or OBS 3.0)

OBS architecture has undergone two generations: OBS 2.0 and OBS 3.0. A newly created bucket is stored in OBS 3.0 by default, and the bucket version is OBS 3.0. However, previously created buckets are still in OBS 2.0.

Basic OBS features and functions are supported by both OBS 3.0 and OBS 2.0. Some new features are supported only by OBS 3.0, such as image processing and cross-region replication.

You can check the bucket version on OBS Console or use the **Head Bucket** API to check whether your bucket is in OBS 2.0 or OBS 3.0. The details are as follows:

### **Method 1: Log in to OBS Console and check the basic bucket information.**

If **Bucket Version** is **3.0**, the bucket is stored in OBS 3.0. If not, the bucket is stored in OBS 2.0.

## Method 2: Use the Head Bucket API to check the bucket version.

### Sample Request:

```
HEAD / HTTP/1.1
Host: bucketname.obs.eu-west-101.myhuaweicloud.eu
Accept: */*
Date: WED, 01 Jul 2015 02:23:25 GMT
Authorization: auth string
```

### Sample Response:

```
HTTP/1.1 200 OK
Server: OBS
x-obs-request-id: BF2600000163D80E4C5F20FDD5BD0085
Content-Type: application/xml
x-obs-version: 3.0
x-obs-id-2: 32AAAQAAEAAABAAAQAAEAAABAAAQAAEAAABCS8wS9l00ll4oMWmdniV7XmdAvfewrQq
Date: WED, 01 Jul 2015 02:23:25 GMT
Content-Length: 0
```

In the response message, **x-obs-version: 3.0** indicates that the bucket is stored in OBS 3.0. If this header does not exist or the value of this header is displayed otherwise, the bucket is stored in OBS 2.0.

For details about the **Head Bucket** API, see [Obtaining Bucket Metadata](#).

# 2 Storage Classes

## Scenarios

OBS provides the following storage classes: Standard, Infrequent Access, and Archive. For the billing about different storage classes, see [Storage Space](#).

These storage classes can meet different needs for storage performance and costs.

- The Standard storage class features low latency and high throughput. It is therefore good for storing frequently (multiple times per month) accessed files or small files (less than 1 MB). Its application scenarios include big data analytics, mobile apps, hot videos, and social apps.
- The Infrequent Access storage class is ideal for storing data that is accessed infrequently (less than 12 times a year) yet has quick response requirements. Its application scenarios include file synchronization, file sharing, and enterprise backup.
- The Archive storage class is ideal for archiving data that is rarely accessed (once a year on average). Its application scenarios include data archiving and long-term data backups. The Archive storage class is secure, durable, and inexpensive, and can be used to replace tape libraries. To keep cost low, it may take hours to restore data from the Archive storage class.

## Comparison of Storage Classes

Compared Item	Standard	Infrequent Access	Archive
Feature	Top-notch performance, high reliability and availability	Reliable, inexpensive storage with real-time access	Long-term, inexpensive storage for Archive data
Application scenarios	Cloud applications, data sharing, content sharing, and hot data storage	Web disk applications, enterprise backup, active archiving, and data monitoring	Storage of archives, medical imaging data, and videos, as well as replacement of tape libraries

Compared Item	Standard	Infrequent Access	Archive
Durability	99.999999999%	99.999999999%	99.999999999%
Availability	99.99%	99%	99%
Minimum measurement unit <sup>a</sup>	64 KB	64 KB	64 KB
Minimum storage duration <sup>b</sup>	N/A	30 days	90 days
Data retrieval	N/A	Billed for each GB retrieved	Data can be restored at a standard or an expedited speed. The time required varies depending on the restore speed. For details, see <a href="#">Restoring Objects from the Archive Storage</a> . Billed for each GB retrieved

 **NOTE**

a: Minimum measurement unit refers to a minimum billable object size. For example, a 32 KB Standard object will be billed as if it were 64 KB.

b: Minimum storage duration refers to a minimum billable storage duration. For example, an Infrequent Access object deleted when it was stored for 20 days will be billed for a full 30 days.

## Bucket Storage Classes and Object Storage Classes

You can specify the storage class for a bucket when creating the bucket. You can also change the storage class of a bucket after the bucket is created.

An object inherits the storage class of the bucket where it is uploaded. You can specify a storage class for an object when uploading it, or you can change the object storage class after the object is uploaded.

Changing the storage class of a bucket does not change the storage classes of existing objects in the bucket, but newly uploaded objects inherit the new storage class by default.

 **NOTE**

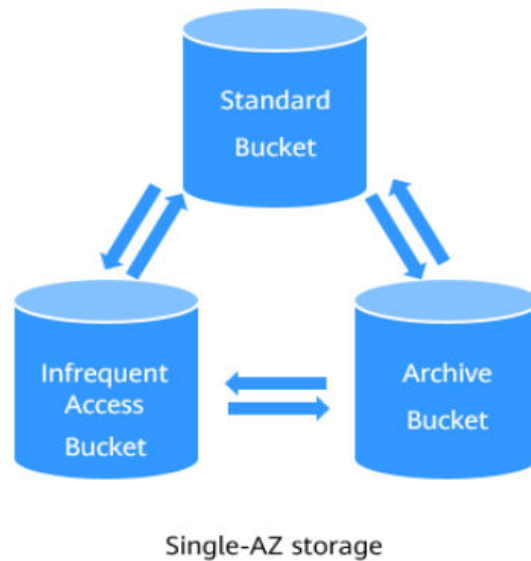
By default, objects in Archive storage class must be restored to Standard storage class before they can be downloaded.

## Changing Bucket Storage Classes

The storage class of a bucket can only be manually changed.

**Figure 2-1** shows how buckets are changed between storage classes. OBS Standard, Infrequent Access, and Archive storage classes support only single-AZ redundancy. The data redundancy policy of a bucket cannot be changed, even when its storage class is changed.

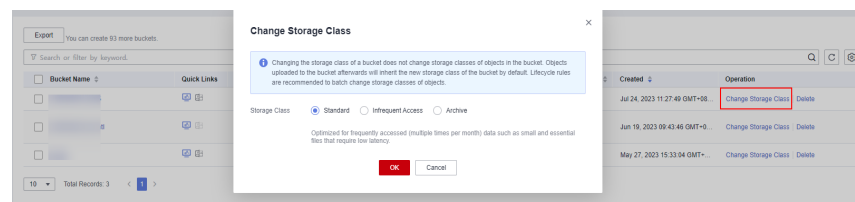
**Figure 2-1** Bucket storage class change rules



**To manually change the storage class of a bucket, do as follows:**

1. Log in to **OBS Console**. In the navigation pane, choose **Object Storage**.
2. In the bucket list, locate the bucket whose storage class you want to change and click **Change Storage Class** in the **Operation** column on the right.
3. Choose a new storage class and click **OK**.

**Figure 2-2** Manually changing the storage class of a bucket



### Precautions:

- Changing the storage class of a bucket does not change the storage classes of existing objects in the bucket. The storage class of an object uploaded later will inherit the new storage class of the bucket by default. You can also **configure lifecycle rules** to change storage classes of objects in a batch.

For example, if **bucket1** is in the Standard storage class and contains **object1**, when **bucket1**'s storage class is transitioned to Infrequent Access, **object1** is

still in the Standard storage class. If you upload **object2** to **bucket1** after the transition, **object2** will be in the Infrequent Access storage class.

- After a bucket is changed from Archive to Standard or Infrequent Access, existing Archive objects in the bucket will not be automatically restored.

## Changing Object Storage Classes

The storage class of an object can be changed manually or automatically. [Figure 2-6](#) shows how objects are transitioned between storage classes.

- Manual: After an object is uploaded, you can manually change its storage class.
  - From Standard to Infrequent Access, or Archive
  - From Infrequent Access to Standard, or Archive
  - From Archive to Standard, or Infrequent Access. Before changing Archive objects, you must restore them first.

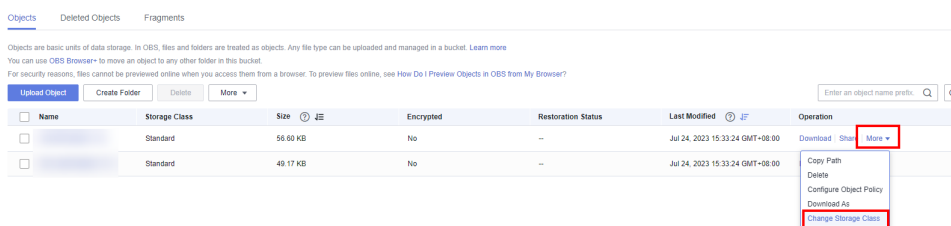
### NOTE

Changing objects from Infrequent Access or Archive to other storage classes incurs restore costs. Select an appropriate change option based on your actual needs.

**To manually change the storage class of an object, do as follows:**

- a. Log in to [OBS Console](#). In the navigation pane, choose **Object Storage**.
- b. In the bucket list, click the bucket name you want. The **Objects** page is displayed.
- c. Change the storage class of objects individually or in a batch.
  - i. Individually: In the object list, locate the desired object and choose **More > Change Storage Class** in the **Operation** column on the right.

**Figure 2-3** Changing object storage classes individually



The screenshot shows the 'Objects' page in the OBS Console. It features a table with columns for Name, Storage Class, Size, Encrypted, Restoration Status, and Last Modified. Two objects are listed: one with a size of 58.60 KB and another with 49.17 KB, both in the Standard storage class. The 'Operation' column for the first object has a 'More' dropdown menu open, which includes options like 'Copy Path', 'Delete', 'Configure Object Policy', 'Download As', and 'Change Storage Class'. The 'Change Storage Class' option is highlighted with a red box.

Name	Storage Class	Size	Encrypted	Restoration Status	Last Modified	Operation
	Standard	58.60 KB	No	--	Jul 24, 2023 15:33:24 GMT+08:00	Download Show More
	Standard	49.17 KB	No	--	Jul 24, 2023 15:33:24 GMT+08:00	Copy Path Delete Configure Object Policy Download As Change Storage Class

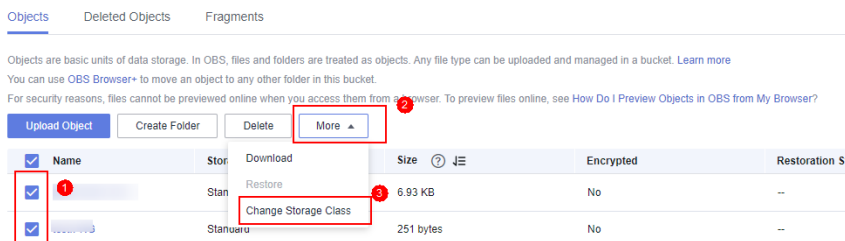
- ii. In a batch: In the object list, select the desired objects and choose **More > Change Storage Class** above the object list.

### NOTE

Changing object storage classes in batches is currently only available in some regions. If you cannot find the button for this batch operation, check whether this operation is supported in your region.

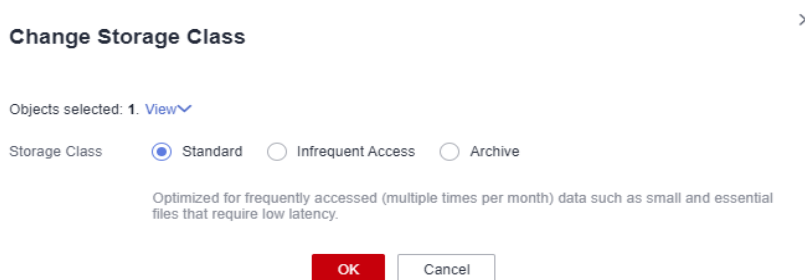


**Figure 2-4** Changing object storage classes in a batch



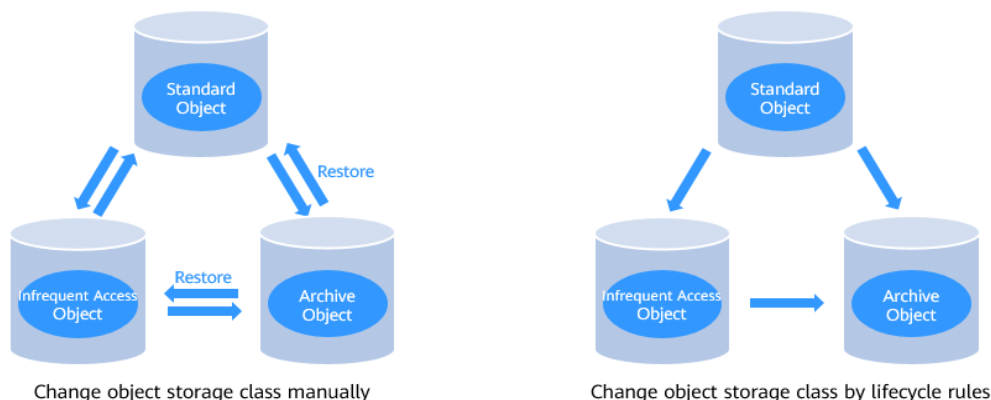
d. Choose a new storage class and click **OK**.

**Figure 2-5** Choosing a new storage class



- Automatic: You can define an OBS lifecycle rule to automatically transition objects from one storage class to another. For details, see [Lifecycle Management](#). Objects can be automatically transitioned as follows:
  - From Standard to Infrequent Access or Archive
  - From Infrequent Access to Archive.

**Figure 2-6** Object storage class transition rules



**Precautions:**

- After objects in the Standard storage class are transitioned to Infrequent Access or Archive storage class, their restoration status is Unrestored.
- The minimum storage duration is 90 days for Archive storage, and 30 days for Infrequent Access storage. If an object is transitioned to another storage class

before it has been stored for the required minimum storage duration, you need to pay for a full 90, or 30 days.

## How to Use

You can use OBS Console or APIs to configure storage classes for buckets and objects.

Tool	Reference
OBS Console	<a href="#">Specifying a Storage Class During Bucket Creation</a> <a href="#">Specifying a Storage Class During Object Uploads</a>
API	<a href="#">Setting the Default Storage Class for a Bucket</a> <a href="#">Specifying a Storage Class During Bucket Creation</a> (adding the <b>x-obs-storage-class</b> header) <a href="#">Specifying a Storage Class During Object Uploads with PUT</a> (adding the <b>x-obs-storage-class</b> header) <a href="#">Specifying a Storage Class During Object Uploads with POST</a> (adding the <b>x-obs-storage-class</b> header)

# 3 Bucket Management

---

## 3.1 Bucket Overview

A bucket is a virtual container used to store **objects** in OBS. OBS offers a flat structure based on buckets and objects. This structure enables all objects to be stored at the same logical layer, rather than being stored hierarchically.

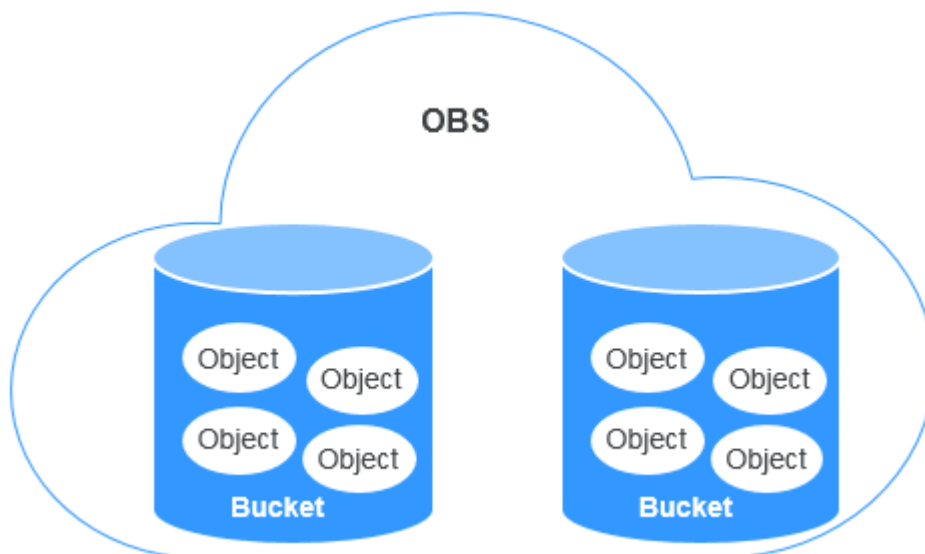
Each bucket has its own properties, such as **a storage class**, access permissions, and **a region**. You can create buckets with different storage classes and access permissions in different regions. You can also configure advanced bucket properties to fit a wide range of storage needs.

On OBS, each bucket name must be unique and cannot be changed. The region where a bucket resides cannot be changed once the bucket is created. When you create a bucket, OBS creates a default access control list (ACL) that grants the authorized user permissions on the bucket. Only authorized users can perform operations such as creating, deleting, viewing, and configuring buckets.

An account (including all IAM users under the account) can create a maximum of 100 buckets. You can leverage the fine-grained permission control capability of OBS to properly plan and use buckets. For example, you can create folders in a bucket based on object prefixes and use **fine-grained permission control** to isolate data between different departments. There is no limit on the number and total size of objects in a bucket.

As OBS is based on a RESTful architecture over HTTP and HTTPS, you can use uniform resource locators (URLs) to locate resources.

**Figure 3-1** illustrates the relationship between buckets and objects in OBS.

**Figure 3-1** Buckets and objects

You can [use different methods](#) to [create buckets](#) based on your use habits and storage needs. After a bucket is created, you can use different ways to [upload files \(data\) to the bucket](#), where these files are stored as objects. In OBS, buckets and objects are located in [different regions](#). You can use different methods to access the same bucket and resources in the same region.

## 3.2 Creating a Bucket

### Scenarios

Buckets are containers that store objects in OBS. You need to create a bucket before you can start storing data in OBS.

### Prerequisites

To create a bucket, make sure you have a registered account, sufficient account balance, access keys (AK and SK), and the endpoint information. For details, see [Getting Started](#).

### Constraints

- After a bucket is created, its name and region cannot be changed. Specify a proper region and bucket name when creating the bucket.
- An account (including all IAM users under the account) can create a maximum of 100 buckets. You can leverage the fine-grained permission control capability of OBS to properly plan and use buckets. For example, you can create folders in a bucket based on object prefixes and use [fine-grained permission control](#) to isolate data between different departments. There is no limit on the number and total size of objects in a bucket.
- A bucket name is part of the access domain name and needs to be resolved. Therefore, the bucket name must conform to the [DNS domain naming rules](#).

When receiving a bucket creation request, the OBS strictly checks the bucket name. A bucket name:

- Must be unique across all accounts and regions. The name of a deleted bucket can be reused for another bucket or a parallel file system at least 30 minutes later after the deletion.
- Must be 3 to 63 characters long. Only lowercase letters, digits, hyphens (-), and periods (.) are allowed.
- Cannot start or end with a period (.) or hyphen (-), and cannot contain two consecutive periods (..) or contain a period (.) and a hyphen (-) adjacent to each other.
- Cannot be formatted as an IP address.

#### NOTICE

If you use a bucket with periods (.) in its name to access OBS, the client will display a message indicating that the bucket is risky, for example, a red alarm may be displayed in the browser security prompt. This is because the SSL wildcard certificate matches only buckets without periods (.) in their names when HTTPS is used for OBS access. We recommend that you avoid using periods (.) in bucket names.

## How to Use

You can use OBS Console, APIs, SDKs, OBS Browser+, or obsutil to create buckets.

Tool	Reference
OBS Console	<a href="#">Creating a Bucket</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Creating a Bucket</a>
OBS Browser+	<a href="#">Creating a Bucket</a>
obsutil	<a href="#">Creating a Bucket</a>

## How to Access a Bucket

After a bucket is created, you can use the domain name to access the bucket. You can assemble the domain name by putting the bucket name and endpoint together, or you can obtain it by [viewing the basic bucket information](#) on OBS Console or OBS Browser+.

An access domain name is structured as follows:

[Structure] *BucketName.Endpoint*

[Example] *bucketname.obs.eu-west-101.dublincloud.cn*

## Causes of Bucket Creation Failures and Solutions

For details, see [Why Am I Unable to Create a Bucket?](#)

## 3.3 Listing Buckets

### Scenarios

You can list all created buckets to view their information.

### How to Use

You can use OBS Console or APIs to list buckets.

Tool	Reference
OBS Console	Log in to OBS Console and select <b>Object Storage</b> . Then all buckets under your account will be displayed.
API	<a href="#">Listing Buckets</a>

## 3.4 Viewing Bucket Information

### Scenarios

After creating a bucket on OBS Console, you can view the bucket information, including its basic information and configurations.

### How to Use

You can use OBS Console or APIs to view basic bucket information.

Tool	Reference
OBS Console	<a href="#">Viewing a Bucket's Basic Information</a>
API	<a href="#">Obtaining Bucket Metadata</a>

## 3.5 Managing Bucket Quotas

### Scenarios

By default, neither the entire OBS system nor any single bucket has limitations on the total size or quantity of objects that can be stored. You can set a quota for a bucket to limit the total size of objects that can be uploaded to the bucket. If the total object size reaches the upper limit, object uploads will fail.

A bucket quota can control object uploading only after the quota is set. If the bucket quota is less than the capacity of the uploaded objects, the existing objects will not be deleted, but new objects cannot be uploaded. In this case, you can upload new objects only after deleting some existing objects until the used storage capacity is less than the quota limit.

 **NOTE**

- A bucket quota must be a non-negative integer, in bytes. The maximum value is  $2^{63} - 1$ .
- OBS does not provide an API for deleting bucket quotas. You can set the bucket quota to 0 to cancel the quota limit.

## How to Use

You can use APIs to manage bucket quotas.

Tool	Reference
API	<a href="#">Configuring Bucket Storage Quota</a>

## 3.6 Deleting Buckets

### Scenarios

You can delete unwanted buckets to free up the quota of buckets. Make sure that a bucket is emptied before you delete it.

An empty bucket must meet the following requirements:

- There is no object or any historical version of an object in the bucket.
- There is not any incomplete multipart upload in the bucket. In other words, there are no fragments in the bucket.

 **NOTE**

- If versioning is enabled for the bucket, ensure that all historical versions and versions with the **Delete Marker** (which are also considered as historical versions) have been deleted.
- The name of a deleted bucket can be reused at least 30 minutes after the deletion.

## How to Use

You can use OBS Console or APIs to delete buckets.

Tool	Reference
OBS Console	<a href="#">Deleting a Bucket</a>
API	<a href="#">Deleting a Bucket</a>

## Causes of Bucket Deletion Failures and Solutions

For details, see [Why Can't I Delete a Bucket?](#)



# 4 Object Management

---

## 4.1 Object Overview

An object is the basic unit of data storage on OBS. It consists of object data and object metadata that describes object attributes. Data uploaded to OBS is stored as objects in **buckets**.

An object consists of data, metadata, and a key.

- A key specifies the name of an object. An object key is a UTF-8 string up to 1,024 characters long. Each object is uniquely identified by a key within a bucket.
- Metadata describes an object, and can be system-defined or user-defined. The metadata is a set of key-value pairs that are assigned to the object stored in OBS.
  - System-defined metadata is automatically assigned by OBS for processing objects. Such metadata includes Date, Content-Length, Last-Modified, Content-MD5, and more.
  - User-defined metadata is specified when you upload objects and is used to describe objects.
- Data refers to the content of the object.

Generally, objects are managed as files. However, OBS is an object-based storage service and there is no concept of files and folders. For easy data management, OBS provides a method to simulate folders. By adding a slash (/) in an object name, for example, **test/123.jpg**, you can simulate **test** as a folder and **123.jpg** as the name of a file under the **test** folder. However, the object key remains **test/123.jpg**.

When uploading an object, you can specify a **storage class** for it. If you do not specify a storage class, the object inherits the storage class of the bucket. After the object is uploaded, you can also change its storage class.

You can **upload files (data) to a bucket** in **the ways** you like based on your habits and service scenarios. OBS then stores the files as objects in the bucket. In OBS, buckets are located in **different regions**. No matter what method you use, you can access the same bucket and its resources in the same region.

## Guidelines on Naming Object Keys

Although any UTF-8 characters can be used in an object key name, naming object keys according to the following guidelines can help maximize the object keys' compatibility with other applications. Ways to analyze special characters vary with applications. The following guidelines help object key names substantially meet the requirements of DNS, web security characters, XML analyzers and other APIs.

The following character sets can be freely used in key names.

Alphanumeric characters (also known as unreserved characters)	[0-9a-zA-Z]
Special characters (also known as reserved characters)	Exclamation mark (!) Hyphen (-) Underscore (_) Period (.) Asterisk (*) Single quotation mark (') Left bracket "(" Right bracket ")"

The following are examples of valid object key names:

```
4my-organization
my.great_photos-2014/jan/myvacation.jpg
videos/2014/birthday/video1.wmv
```

## Percent-Encoding of Reserved Characters

If a reserved character has a special meaning (known as reserved purpose) in a URI and the character must be used for other purposes in the URI, this character must be percent-encoded. Use UTF-8 to encode non-ASCII characters. Otherwise, the names of the objects that are uploaded to OBS may be different from what is expected. For example, if reserved character "/" is used as the delimiter of path components in a URI, the character has a special meaning (separating a bucket name from an object name). If "/" is used in a component of the path in a URI, use three characters "%2F" or "%2f" to replace "/". The reserved character " " must be encoded as "%20". For example, the string "abc d" will be encoded as "abc%20d".

## Characters That May Require Special Processing

Characters that require encoding in a key name

- Ampersand (&)
- Dollar sign (\$)
- Semicolon (;)
- Colon (:)

- Plus sign (+): OBS decodes plus signs (+) in a request URI into spaces. If an original object key name contains plus signs (+), it must be encoded into %2B before being put into the request URI.
- Space: A large number of consecutive spaces may be lost in some cases.
- Equality sign (=)
- At sign (@)
- Comma (,)
- Question mark (?)
- ASCII characters: 00–1F in hexadecimal form (0–31 in decimal form) and 7F (127 in decimal form)

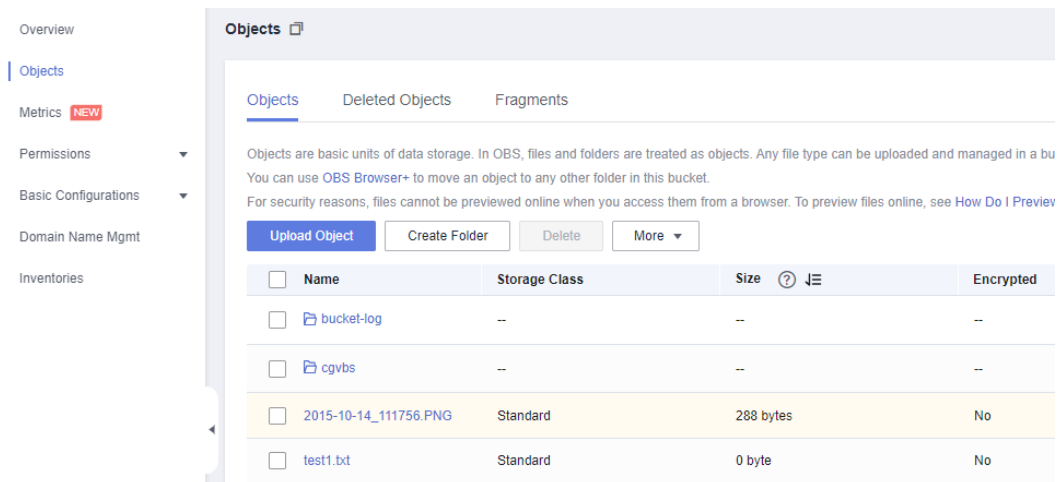
Avoid using the following characters in key names, because these characters require a lot of special processing to keep consistency across all applications.

- Backslash (\)
- Left brace {}
- Non-printable ASCII characters (128–255 decimal characters)
- Insert symbol (^)
- Right brace {}
- Percentage character (%)
- Accent/Untick (`)
- Right square bracket (])
- Quotation mark
- Greater than sign (>)
- Left square bracket ([)
- Tilde (~)
- Less than sign (<)
- Number sign (#)
- Vertical bar (|)

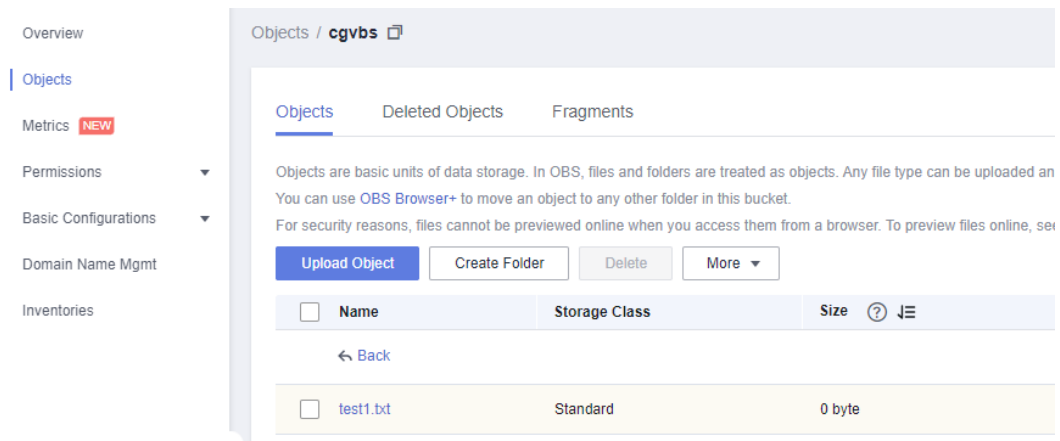
Note that OBS adopts a flat structure, where you create buckets and store objects in buckets. There are no sub-buckets or sub-folders in the structure. However, you can use key name prefixes and delimiters to deduce the logical structure like OBS Console does. The folder concept is available on OBS Console. Assume that your bucket (**companybucket**) contains four objects with the following object keys:

- bucket-log/log01.txt
- cgvs/test1.txt
- 2015-10-14\_111756.png
- test1.txt

OBS Console uses a key name prefix (**bucket-log/** or **cgvs/**) and separator (/) to display the folder structure, as shown in the following figure.



The **2015-10-14\_111756.png** and **test1.txt** keys do not have a prefix, so they appear at the root level of the bucket. If you open the **cgvbs/** folder, you will see that it contains the **test1.txt** object.



Assume that your bucket (**companybucket**) contains two objects with the following object keys:

- obj
- 1/./obj

If you call an API (for example, using the SDK) to obtain these two objects, you will successfully get them. If you use OBS Console to access them, you will actually get object **obj** based on the relative relationship. Since **./** has special semantics in URIs, avoid using it in object keys.

Object names starting with consecutive periods (**..**) are not allowed.

## 4.2 Uploading Objects

## 4.2.1 Uploading an Object

### Scenarios

You can upload files or folders to an existing OBS bucket. If you want to classify files to be uploaded, create a folder and upload related files to the folder.

For details about how to create a folder, see [Creating a Folder](#).

These files can be texts, images, videos, or any other type of files.

---

#### NOTICE

OBS allows you to upload objects to buckets in a specified region, but Huawei Cloud does not detect the object content you uploaded. If your object uploads involve cross-border transfer, ensure that your use complies with relevant laws and regulations.

---

### Constraints

#### Size and number of objects to upload

- OBS Console puts limits on the size and number of files you can upload.
  - In regions that support batch uploads, a maximum of 100 files can be uploaded at a time, with a total size of no more than 5 GB. If you upload only one file in a batch, this file cannot exceed 5 GB in size.
  - In regions that do not support batch uploads, only one file can be uploaded at a time, with a size of no more than 50 MB.
- With OBS Browser+ or obsutil, you can upload files smaller than 48.8 TB. OBS Browser+ allows you to upload a maximum of 500 files at a time. There is no limit on the number of files you can upload using obsutil at a time.
- With PUT, POST, or append methods of the OBS SDKs or API, you can upload files smaller than 5 GB.
- With multipart uploads of the OBS SDKs or API or resumable uploads of the OBS SDKs, you can upload files smaller than 48.8 TB.
- You can specify a quota for a bucket to limit the bucket's capacity for storing the uploaded objects. For details, see [Managing Bucket Quotas](#).

#### Object naming

- See [Guidelines on Naming Object Keys](#).

#### Batch operations

#### Folder uploads

- You can upload folders using OBS Console, OBS Browser+, or obsutil.

#### Others

- If versioning is disabled and the name of a newly uploaded file is the same as that of a file in the bucket, the newly uploaded file automatically overwrites the existing file and does not retain the ACL information of the existing file. If

the name of the newly uploaded folder is the same as that of a folder in the bucket, the two folders will be merged, and files in the new folder will overwrite namesake files in the old folder.

- If versioning is enabled and the name of a newly uploaded file is the same as that of a file in the bucket, a new version is added to the existing file. For details about versioning, see [Versioning](#).

## Upload Methods

Upload Method	Description
PUT (Streaming upload)	Use the PUT or POST method when the size of the file to be uploaded is less than 5 GB. For details about the differences between the two upload methods, see <a href="#">What Are the Differences Between PUT and POST Upload Methods?</a>
POST (browser-based upload)	
Multipart upload	Use this method when the size of the file to be uploaded is greater than 5 GB and less than 48.8 TB. For details, see <a href="#">Uploading Objects Using a Multipart Upload</a> .
Resumable transfer	Uploading large files often fails due to poor network conditions or program breakdowns. It is a waste of resources to restart the upload process upon an upload failure, and the restarted upload process may still suffer from the unstable network. To resolve such issues, you can use the API for resumable upload, whose working principle is to divide the to-be-uploaded file into multiple parts and upload them separately. This method saves resources and improves efficiency upon the re-upload, and speeds up the upload process by concurrently uploading parts. For details, see <a href="#">Does OBS Support Resumable Data Transfer?</a>
Synchronous upload of incremental objects	This method synchronizes all content in the local source path to the specified target bucket in OBS, ensuring that the content is consistent between the local path and the target bucket. Incremental synchronization has the following meanings: 1) Increment: Compare the source file with the target object and upload only the source file that has changes. 2) Synchronization: After the command is executed, ensure that the local source path is a subset of the target bucket specified by OBS. That is, any file in the local source path has its corresponding object in the target bucket on OBS.
Appendable upload	The AppendObject operation adds data to the end of an object in a specified bucket. If there is no namesake object in the bucket, a new object is created. For details, see <a href="#">Appending an Object</a> .

## How to Use

You can use OBS Console or APIs to upload objects.

Tool	Reference
OBS Console	<a href="#">Uploading a File</a>
API	<a href="#">Operations on Objects</a>

## Related Operations

When uploading an object, you can specify a [storage class](#) for it. After the object is uploaded, you can also change its storage class.

- You can manually change objects among the Standard, Infrequent Access, Archive storage classes. Objects in the Archive storage class must be restored before they can be manually changed to the Standard, or Infrequent Access storage class. Changing the storage class of objects in the Infrequent Access or Archive storage class involves retrieval costs, so you are advised to change their storage classes based on the access frequency and scenario.
- The minimum storage duration is 90 days for Archive storage, and 30 days for Infrequent Access storage. If an object is transitioned to another storage class before it has been stored for the required minimum storage duration, you need to pay for the remaining days.
- You can also configure a lifecycle rule to transition the storage class of an object. For details, see [Lifecycle Management](#).

## Causes of Object Upload Failures and Solutions

For details, see [Why Am I Unable to Upload an Object?](#)

### 4.2.2 Uploading Objects Using a Multipart Upload

Multipart upload allows uploading a single object as a group of parts separately. Each part is a part of consecutive object data. You can upload object parts in any sequence or independently upload them. A part can be reloaded after an uploading failure, without affecting other parts. After all parts are uploaded, OBS merges these parts to create the object. Generally, if the size of an object reaches 100 MB, multipart upload is recommended. For example, if you want to upload a 500 MB object to an OBS bucket, you can use OBS Browser+ to upload the object using a multipart upload. OBS Browser+ divides the object into small parts and then uploads the parts. Alternatively, you can call the multipart upload API, improving upload efficiency and reducing failures.

Multipart upload provides the following benefits:

- Improving throughput: You can upload parts in parallel to improve throughput.
- Quick recovery from any network failures: Small-size parts can minimize the impact of failed uploading caused by network errors.

- Convenient suspension and resuming of object uploading: You can upload parts at any time. A multipart upload does not have a validity period. You must explicitly complete or cancel the multipart upload.
- Starting uploading before knowing the size of an object: You can upload an object while creating it.

The multipart upload API allows uploading a large object in multiple parts. You can upload a new large object or create a copy of an existing object using this API.

The procedure for uploading multiple sections is as follows: Starting uploading (initializing the upload task), uploading parts, and completing uploading (merging the uploaded parts). Upon receiving a part merging request, OBS merges the uploaded parts to create a new object. The object can be accessed like other objects.

You can list all the ongoing multipart upload tasks or obtain the list of uploaded parts of a specified multipart upload task. The following describes the detailed operations.

## Initiating a Multipart Upload

When you send a request to start multipart upload, OBS returns a response with the upload ID, which is the unique identifier of the multipart upload. This ID must be included in the request for uploading parts, listing uploaded parts, completing a multipart upload, or canceling a multipart upload.

## Uploading a Part

When uploading parts, you must specify the upload ID and part numbers. You can select any part number between 1 and 10,000. A part number uniquely identifies a part and its location in the object you are uploading. If the number of an uploaded part is used to upload a new part, the uploaded part will be overwritten. Whenever you upload a part, OBS returns the ETag header in the response. For each part upload task, you must record the part numbers and ETag values. These part numbers and ETag values are required in subsequent operations of completing the multipart upload task.

---

### NOTICE

After the multipart upload task is initialized and one or more parts are uploaded, you must merge the parts or cancel the multipart upload task. Otherwise, you have to pay for the storage fee of the uploaded parts. OBS releases the storage and stops charging the storage fee only after the uploaded parts are merged or the multipart upload task is canceled.

---

When multiple concurrent upload operations are performed for the same part of an object, the server complies with the Last Write Win policy, but the time referred in Last Write is the time when the part metadata is created. To ensure data accuracy, the client must be locked during the concurrent upload for the same part of an object. Concurrent upload for different parts of an object does not require the client to be locked.



## Copying a Part

After creating a multipart upload job, you can specify upload IDs and upload parts for the specified upload task. You can also call the API for part copying to add parts. A part of an object or the whole object can be copied as a part.

---

### NOTICE

You cannot determine whether a request is successful only based on the **status\_code** in the returned HTTP header. If **200** is returned for **status\_code**, the server has received the request and started to process the request. The copy is successful only when the body in the response contains ETag.

---

If you copy the source object as a part called part1 and another part1 already exists before the copy operation, the original part1 will be overwritten by the new one after the copy operation. After the copy succeeds, only the new part1 is displayed. Data of the old part1 will be deleted. Therefore, ensure that the target part does not exist or has no value when copying a part. Otherwise, data may be deleted by mistake. The source object does not change during the copying.

## Merging Parts and Canceling a Multipart Upload Task

When merging parts, OBS creates an object by standardizing multiple parts in ascending order. If any object metadata is provided in the initialization of a part upload task, OBS associates the metadata with the object. After the multipart upload is complete, the parts will no longer exist. A part merging request must contain the upload ID, part numbers, and a list of corresponding ETag values. OBS responses include the ETag that uniquely identifies composite object data. The ETag is not the MD5 hash value of the object data. You can cancel a multipart upload task. After a multipart upload task is canceled, the upload ID cannot be used to upload any part. Then, OBS releases the storage of all uploaded parts. If you stop an ongoing multipart upload, the uploading will still complete (the result can be successful or failed). To release the storage capacity occupied by all uploaded parts, cancel the multipart upload after the entire task is complete.

---

### NOTICE

If 10 parts are uploaded but only nine parts are selected for merge, the parts that are not merged will be automatically deleted by the system. The parts that are not merged cannot be restored after being deleted. Before merging the parts, adopt the API used to list the parts that have been uploaded to check all parts to ensure that no part is missed.

## Listing Uploaded Parts

You can list the parts of a specified multipart upload task or the parts of all the multipart upload tasks in progress. Information about uploaded parts in a specific multipart upload will be returned for a request to list uploaded parts. For each request to list uploaded parts, OBS returns information about uploaded parts in the specific multipart upload. Information about a maximum of 1000 parts can be

returned. If more than 1000 parts are uploaded in a multipart upload, you need to send multiple requests to list all uploaded parts. The list of uploaded parts does not include merged parts.

**NOTICE**

A returned list can only be used for verification. After a multipart upload is complete, the result in the list is no longer valid. However, when part numbers and the ETag values returned by OBS are uploaded, the list of part numbers specified by the user will be reserved.

## Listing Multipart Uploads

You can obtain the list of initialized multipart upload tasks by listing the multipart upload tasks in the bucket. Initialized multipart upload tasks refer to the multipart upload tasks that are not merged or canceled after initialization. A maximum of 1000 multipart upload tasks can be returned for each request. If the number of ongoing multipart upload tasks exceeds 1000, you need to send more requests to query the remaining tasks.

**Table 4-1** lists the restrictions on listing multipart uploads.

**Table 4-1** Restrictions on listing multipart uploads

Item	Restriction
Object size	Up to 48.8 TB
Maximum number of parts for each upload task	10,000
Part number	1–10,000 (included)
Part size	The part size is between 5 MB to 5 GB. The size of the last part is between 0 bytes to 5 GB.
Maximum number of uploaded parts returned in response to the request for listing uploaded parts.	1000
Maximum number of initialized multipart upload tasks returned in response to the request for listing initialized multipart tasks.	1000

## Multipart Upload Operations and Permissions

You can perform multipart upload only after being granted with the permission. You can use ACLs, bucket policies, or user policies to grant users the permission. The following table lists multipart upload operations and the required permissions that can be granted by ACLs, bucket policies, or user policies.

Operation	Required Permission
Initiating a multipart upload	<p>To perform this operation, you need to have the PutObject permission.</p> <p>A bucket owner can allow trustees to perform the PutObject operation.</p>
Uploading parts	<p>To perform this operation, you need to have the PutObject permission.</p> <p>Only the initiator of a multipart upload can upload parts. The bucket owner must grant the multipart upload initiator the PutObject permission so that the initiator can upload parts of the object.</p>
Copying parts	<p>To perform this operation, you need to have the PutObject permission as well as the GetObject permission on the object to be copied.</p> <p>Only the initiator of a multipart upload can copy parts. The bucket owner must grant the multipart upload initiator the PutObject permission so that the initiator can upload parts of the object.</p>
Assembling parts	<p>To perform this operation, you need to have the PutObject permission.</p> <p>Only the initiator of a multipart upload can assemble parts. The bucket owner must grant the multipart upload initiator the PutObject permission so that the initiator can complete the multipart upload.</p>
Canceling a multipart upload	<p>To perform this operation, you need to have the AbortMultipartUpload permission.</p> <p>By default, only the bucket owner and multipart upload initiator have this permission. In addition to the default configuration, the bucket owner can allow trustees to perform this operation. The bucket owner can also deny any trustees performing this operation.</p>

Operation	Required Permission
Listing uploaded parts	<p>To perform this operation, you need to have the ListMultipartUploadParts permission.</p> <p>By default, the bucket owner can list the uploaded parts of any multipart upload to the bucket. The multipart upload initiator can list the uploaded parts of a specific multipart upload.</p> <p>In addition to the default configuration, the bucket owner can allow trustees to perform this operation. The bucket owner can also deny any trustees performing this operation.</p>
Listing multipart uploads	<p>To list multipart upload tasks to the bucket, you need to have the ListBucketMultipartUploads permission.</p> <p>In addition to the default configuration, the bucket owner can allow trustees to perform this operation.</p>

## REST APIs Applicable to Multipart Upload

The following sections in the *Object Storage Service API Reference* describe REST APIs relevant to multipart upload.

- Listing Initiated Multipart Uploads in a Bucket
- Initiating a Multipart Upload
- Multipart Upload
- Uploading a Part of an Object - Copy
- Listing Uploaded Parts of an Object
- Completing a Multipart Upload
- Canceling a Multipart Upload Task

## 4.3 Downloading Objects

### Scenarios

You can download objects from OBS to a local computer.

OBS supports batch download of objects. For details, see [Does OBS Support Batch Download?](#)

## Constraints

- Objects in the Archive storage class can be downloaded only when they are in the **Restored** state.

## How to Use

You can use OBS Console, OBS Browser+, APIs, SDKs, or obsutil to download objects and use OBS Browser+ and obsutil to download folders.

Tool	Reference
OBS Console	<a href="#">Downloading a File</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Downloading Objects</a>
OBS Browser+	<a href="#">Downloading a File or Folder</a>
obsutil	<a href="#">Downloading an Object</a>

## 4.4 Managing Folders

### 4.4.1 Creating a Folder

#### Scenarios

You can create a folder in a bucket to facilitate classification and management of data stored in OBS.

Unlike a file system, OBS does not involve the concepts of file and folder. For easy data management, OBS provides a method to simulate folders. On OBS Console or OBS Browser+, an object is simulated as a folder by adding a slash (/) to the end of the object name. When objects are listed by calling the API, the obtained object name is the path of the object, and the content following the last slash (/) is the actual object name. If the path ends with a slash (/), it indicates that the object is a folder. The hierarchical depth of the object does not affect the performance of accessing the object.

#### How to Use

You can use OBS Console, APIs, SDKs, OBS Browser+, or obsutil to create a folder.

Tool	Reference
OBS Console	<a href="#">Creating a Folder</a>

Tool	Reference
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page. <b>NOTE</b> To create a folder in OBS is to create an object whose size is 0 and whose name ends with a slash (/), in essential.
API	- <b>NOTE</b> To create a folder in OBS is to create an object whose size is 0 and whose name ends with a slash (/), in essential.
OBS Browser+	<a href="#">Creating a Folder</a>
obsutil	<a href="#">Creating a Folder</a>

## 4.5 Other Object Operations

### 4.5.1 Listing Objects

#### Scenarios

You can view the list of created objects by listing objects.

#### How to Use

You can use OBS Console, APIs, SDKs, or obsutil to list objects.

Tool	Reference
OBS Console	Log in to OBS Console and choose <b>Object Storage</b> in the navigation pane. On the displayed page, click a bucket name to go to the <b>Objects</b> page and view all objects in the bucket.
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Listing Objects in a Bucket</a>
obsutil	<a href="#">Listing Objects</a>

## 4.5.2 Copying Objects

### Scenarios

You can copy objects in OBS. An object of up to 5 GB can be copied by each operation. To copy an object larger than 5 GB, use the multipart upload API. Specifically, this copy operation allows you to:

- Create a copy for an object.
- Rename an object by creating a copy for it and deleting the source object.
- Edit object metadata. Each object has metadata, which is a set of name-value pairs. You can set object metadata when uploading the object. After you upload the object, you can modify its metadata using an API. For details, see [Modifying Object Metadata](#). You can also create an object copy and set the metadata. In the copy operation, set the target object to be the same as the source object.

Each object contains metadata, including system metadata and user-defined metadata. You can control some system metadata. When you copy an object, user-controlled system metadata and user-defined metadata are copied too. OBS resets the metadata controlled by the system. For example, when you copy an object, the OBS resets the creation date of the copied object. In the copy request, you do not need to set such value.

When copying an object, you may want to update some metadata. For example, if the source object is configured to be in Standard storage, you may want to change the storage class of the object copy to Infrequent Access. You may also want to modify some user-defined metadata of the source object. If you want to modify metadata, even only one piece of metadata, that can be configured by users (defined by users or the system), specify all the metadata that can be configured by users on the source object in the request.

#### NOTE

- OBS allows you to replicate the service data stored in OBS to a specified region, but Huawei Cloud does not detect the stored data and is not responsible for the legal compliance of your use of OBS. If your replication involves cross-border transfer, ensure that your use complies with relevant laws and regulations.
- When versioning is disabled for a bucket, if you make a copy of **objecta** and save it as **objectb**, and an object named **objectb** already exists before this copy, the new **objectb** will overwrite the existing one. After the copy succeeds, only new **objectb** can be downloaded because old **objectb** has been deleted. Therefore, before copying an object, make sure that there is no object with the same name as the object copy to prevent data from being deleted mistakenly. Source object **objecta** does not change during the copying.
- You cannot determine whether a request is executed successfully only using **status\_code** in the header returned by HTTP. If 200 in **status\_code** is returned, the server has received the request and starts to process the request. The body in the response shows whether the request is executed successfully. The request is executed successfully only when the body contains ETag; otherwise, the request fails to be executed.

### How to Use

You can use APIs, SDKs, OBS Browser+, or obsutil to copy objects.

Tool	Reference
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Copying Objects</a>
OBS Browser+	<a href="#">Copying a File or Folder</a>
obsutil	<a href="#">Copying Objects</a>

## 4.5.3 Viewing Object Information

### Scenarios

After uploading an object, you can view the information or properties of the object.

### How to Use

You can use OBS Console, APIs, SDKs, or obsutil to view the information of an object.

Tool	Reference
OBS Console	-
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Querying Object Metadata</a>
obsutil	<a href="#">Querying Object Properties</a>

## 4.5.4 Sharing Objects

### Scenarios

You can share files with other users temporarily or permanently.

- Temporary sharing:  
Share a file to other users through the temporary URL of the file. All shared URLs are temporary with a validity period.

A temporary URL consists of the access domain name and the temporary authentication information of a file. Example:

```
https://bucketname.obs.eu-west-101.dublincloud.cn:443/image.png?AccessKeyId=xxx&Expires=xxx&x-obs-security-token=xxx&Signature=xxx
```



The temporary authentication information contains the **AccessKeyId**, **Expires**, **x-obs-security-token**, and **Signature** parameters. The **AccessKeyId**, **x-obs-security-token**, and **Signature** parameters are used for authentication. The **Expires** parameter specifies the validity period of the authentication.

For details about the temporary authentication method and parameters, see [Authentication of Signature in a URL](#) in the *Object Storage Service API Reference*.

- Permanent sharing:

If you set the permission for an object to allow anonymous users to read it, anonymous users can access the object through the URL that you shared. For details about how to configure permissions, see [Granting Public Read Permissions on Objects to Anonymous Users](#).

The method of using a browser to access objects varies depending on the object type. You can directly open **.txt** and **.html** files using a browser. However, when you open **.exe** and **.dat** files using a browser, the files are automatically downloaded to your local computer.

For details about how to obtain the object access URL, see [How Do I Obtain the Access Path to an Object?](#)

## Constraints

- Objects in the Archive storage class can be shared only after they have been restored.
- An authorization code is not required for temporarily sharing a file.
- Only buckets of version 3.0 or later support temporary file sharing. For details about how to query the bucket version, see [Checking OBS Version \(OBS 2.0 or OBS 3.0\)](#).
- The following table describes the URL validity period of the files that are temporarily shared using different tools.

Tool	URL Validity Period
OBS Console	The URL is valid for 1 minute to 18 hours. After an object is shared, the system will generate a URL that contains the temporary authentication information, valid for five minutes since its generation by default. Each time you change the validity period of a URL, OBS obtains the authentication information again to generate a new URL for sharing. The new URL becomes valid at the moment the validity period is changed.
API	Configure parameter <b>Expires</b> to specify when a temporary authorization expires. The temporary authorization expires in 24 hours.

## How to Use

You can use OBS Console, APIs, or obsutil to share files.

Tool	Reference
OBS Console	<a href="#">Sharing a File</a>
API	<a href="#">Authentication of Signature in a URL</a>
obsutil	<a href="#">Generating the Download Link of an Object</a>

## Related Operations

To download a generated temporary URL with wget, use "" and -O to specify the file name, or an error will occur.

Example:

```
[root@ecs-test ~]# wget "Temporary URL" -O abc
```

In the preceding command, *abc* indicates a new file name.

## 4.5.5 Managing Object Metadata

### Scenarios

Object metadata is a set of name-value pairs that describe the object, and are used for object management.

Currently, only the metadata defined by the system is supported.

The metadata defined by the system is classified into the following types: system-controlled and user-controlled. For example, metadata such as **Last-Modified** is controlled by the system and cannot be modified. You can modify the metadata such as **ContentLanguage** through the API. The metadata that can be modified is described as follows:

Table 4-2 OBS metadata

Name	Description
ContentDisposition	<p>Provides a default file name for the object that is being requested. When an object is being downloaded or accessed, the file with the default file name is directly displayed in the browser or a download dialog box is displayed if the file is being accessed.</p> <p>For example, select <b>ContentDisposition</b> as the metadata name and enter <b>attachment;filename="testfile.xls"</b> as the metadata value for an object. If you access the object through a link, a dialog box is directly displayed for downloading objects, and the object name is changed to <b>testfile.xls</b>.</p> <p>For details, see the definition about ContentDisposition in HTTP.</p>
ContentLanguage	<p>Indicates the language or languages intended for the audience. Therefore, a user can differentiate according to the user's preferred language. For details, see the definition about ContentLanguage in HTTP.</p>
WebsiteRedirectLocation	<p>Redirects an object to another object or an external URL. The redirection function is implemented using static website hosting.</p> <p>For example, you can perform the following operations to implement object redirection:</p> <ol style="list-style-type: none"><li>1. Set metadata of object <b>testobject.html</b> in the root directory of bucket <b>testbucket</b>. Select <b>WebsiteRedirectLocation</b> for <b>Name</b> and enter <b>http://www.example.com</b> for <b>Value</b>.</li></ol> <p><b>NOTE</b> OBS only supports redirection for objects in the root directory of a bucket.</p> <ol style="list-style-type: none"><li>2. Configure static website hosting for bucket <b>testbucket</b>, and set the object <b>testobject.html</b> in the bucket as the default home page of the hosted static website.</li><li>3. If you access object <b>testobject.html</b> through the URL link provided on the <b>Configure Static Website Hosting</b> page, the access request is redirected to <b>http://www.example.com</b>.</li></ol>

Name	Description
ContentEncoding	Content encoding format when an object is downloaded. The options are as follows: <ul style="list-style-type: none"><li>• Standard: <b>compress</b>, <b>deflate</b>, <b>exi</b>, <b>identity</b>, <b>gzip</b>, and <b>pack200-gzip</b></li><li>• Others: <b>br</b>, <b>bzip2</b>, <b>lzma</b>, <b>peerdist</b>, <b>sdch</b>, <b>xpress</b>, and <b>xz</b></li></ul>
CacheControl	Cache behavior of the web page when the specified object is downloaded. <ul style="list-style-type: none"><li>• Cacheability: <b>public</b>, <b>private</b>, <b>no-cache</b>, and <b>only-if-cached</b></li><li>• Expiration time: <b>max-age=&lt;seconds&gt;</b>, <b>s-maxage=&lt;seconds&gt;</b>, <b>max-stale[=&lt;seconds&gt;]</b>, <b>min-fresh=&lt;seconds&gt;</b>, <b>stale-while-revalidate=&lt;seconds&gt;</b>, <b>stale-if-error=&lt;seconds&gt;</b></li><li>• Re-verification and reloading: <b>must-revalidate</b>, <b>proxy-revalidate</b>, and <b>immutable</b></li><li>• Others: <b>no-store</b> and <b>no-transform</b></li></ul>
Expires	Cache expiration time (GMT)
ContentType	File type of an object. For details, see <a href="#">Object Metadata Content-Type</a> .

 NOTE

- When versioning is enabled for a bucket, you can set metadata for objects which are **Latest Version**, but cannot set metadata for objects which are **Historical Version**.
- You cannot configure object metadata for an Archive object.

## Object Metadata Content-Type

When an object is uploaded to OBS through OBS Console or a tool, the system automatically matches the value of **Content-Type** based on the file name extension of the object. When you access an object through a web browser, the system specifies an application to open the object according to the value of **Content-Type**. You can modify the **Content-Type** of an object based on its file name extension.

 NOTE

If you upload an object by calling an API, specify the value of **Content-Type** because the system does not automatically match the value of **Content-Type**.

**Table 4-3** Common Content-Type values

File Name Extension	Content-Type	File Name Extension	Content-Type
.* (binary stream, unknown file type)	application/octet-stream	.7z	application/x-7z-compressed
.001	application/x-001	.301	application/x-301
.323	text/h323	.906	application/x-906
.907	drawing/907	.a11	application/x-a11
.acp	audio/x-mei-aac	.ai	application/postscript
.aif	audio/aiff	.aifc	audio/aiff
.aiff	audio/aiff	.anv	application/x-anv
.asa	text/asa	.asf	video/x-ms-asf
.asp	text/asp	.asx	video/x-ms-asf
.au	audio/basic	.avi	video/avi
.awf	application/vnd.adobe.workflow	.biz	text/xml
.bmp	application/x-bmp	.bot	application/x-bot
.c4t	application/x-c4t	.c90	application/x-c90
.cal	application/x-cals	.cat	application/vnd.ms-pki.seccat
.cdf	application/x-netcdf	.cdr	application/x-cdr
.cel	application/x-cel	.cer	application/x-x509-ca-cert
.cg4	application/x-g4	.cgm	application/x-cgm
.cit	application/x-cit	.class	java/*
.cml	text/xml	.cmp	application/x-cmp
.cmx	application/x-cmx	.cot	application/x-cot
.crl	application/pkix-crl	.crt	application/x-x509-ca-cert
.csi	application/x-csi	.css	text/css
.cut	application/x-cut	.dbf	application/x-dbf

File Name Extension	Content-Type	File Name Extension	Content-Type
.dbm	application/x-dbm	.dbx	application/x-dbx
.dcd	text/xml	.dcx	application/x-dcx
.der	application/x-x509-ca-cert	.dgn	application/x-dgn
.dib	application/x-dib	.dll	application/x-msdownload
.doc	application/msword	.docx	application/vnd.openxmlformats-officedocument.wordprocessingml.document
.drw	application/x-drw	.dot	application/msword
.dwf	Model/vnd.dwf	.dtd	text/xml
.dwg	application/x-dwg	.dwf	application/x-dwf
.dxf	application/x-dxf	.dxb	application/x-dxb
.emf	application/x-emf	.edn	application/vnd.adobe.edn
.ent	text/xml	.eml	message/rfc822
.eps	application/x-ps	.epi	application/x-epi
.etd	application/x-ebx	.eps	application/postscript
.fax	image/fax	.exe	application/x-msdownload
.fif	application/fractals	.fdf	application/vnd.fdf
.frm	application/x-frm	.fo	text/xml
.gbr	application/x-gbr	.g4	application/x-g4
.gif	image/gif	.	application/x-
.gp4	application/x-gp4	.gl2	application/x-gl2
.hmr	application/x-hmr	.hgl	application/x-hgl
.hpl	application/x-hpl	.hpg	application/x-hpgl
.hrf	application/x-hrf	.hqx	application/mac-binhex40

File Name Extension	Content-Type	File Name Extension	Content-Type
.htc	text/x-component	.hta	application/hta
.html	text/html	.htm	text/html
.htx	text/html	.htt	text/webviewhtml
.ico	image/x-icon	.icb	application/x-icb
.iff	application/x-iff	.ico	application/x-ico
.igs	application/x-igs	.ig4	application/x-g4
.img	application/x-img	.iii	application/x-iphone
.isp	application/x-internet-signup	.ins	application/x-internet-signup
.java	java/*	.IVF	video/x-ivf
.jpe	image/jpeg	.jfif	image/jpeg
.jpeg	image/jpeg	.jpe	application/x-jpe
.jpg	application/x-jpg	.jpg	image/jpeg
.jsp	text/html	.js	text/javascript
.lar	application/x-laplayer-reg	.la1	audio/x-liquid-file
.lavs	audio/x-liquid-secure	.latex	application/x-latex
.lmsff	audio/x-la-lms	.lbm	application/x-lbm
.ltr	application/x-ltr	.ls	application/x-javascript
.m2v	video/x-mpeg	.m1v	video/x-mpeg
.m4e	video/mpeg4	.m3u	audio/mpegurl
.man	application/x-troff-man	.mac	application/x-mac
.mdb	application/msaccess	.math	text/xml
.mfp	application/x-shockwave-flash	.mdb	application/x-mdb
.mhtml	message/rfc822	.mht	message/rfc822
.mid	audio/mid	.mi	application/x-mi
.mil	application/x-mil	.midi	audio/mid

File Name Extension	Content-Type	File Name Extension	Content-Type
.mnd	audio/x-musicnet-download	.mml	text/xml
.mocha	application/x-javascript	.mns	audio/x-musicnet-stream
.mp1	audio/mp1	.movie	video/x-sgi-movie
.mp2v	video/mpeg	.mp2	audio/mp2
.mp4	video/mp4	.mp3	audio/mp3
.mpd	application/vnd.ms-project	.mpa	video/x-mpg
.mpeg	video/mpg	.mpe	video/x-mpeg
.mpga	audio/rn-mpeg	.mpg	video/mpg
.mps	video/x-mpeg	.mpp	application/vnd.ms-project
.mpv	video/mpg	.mpt	application/vnd.ms-project
.mpw	application/vnd.ms-project	.mpv2	video/mpeg
.mtx	text/xml	.mpx	application/vnd.ms-project
.net	image/pnetvue	.mxx	application/x-mmx
.nws	message/rfc822	.nrf	application/x-nrf
.out	application/x-out	.odc	text/x-ms-odc
.p12	application/x-pkcs12	.p10	application/pkcs10
.p7c	application/pkcs7-mime	.p7b	application/x-pkcs7-certificates
.p7r	application/x-pkcs7-certreqresp	.p7m	application/pkcs7-mime
.pc5	application/x-pc5	.p7s	application/pkcs7-signature
.pcl	application/x-pcl	.pci	application/x-pci
.pdf	application/pdf	.pcx	application/x-pcx
.pdx	application/vnd.adobe.pdx	.pfx	application/x-pkcs12



File Name Extension	Content-Type	File Name Extension	Content-Type
.pgl	application/x-pgl	.pic	application/x-pic
.pko	application/ vnd.ms-pki.pko	.pl	application/x-perl
.plg	text/html	.pls	audio/scpls
.plt	application/x-plt	.png	image/png
.png	application/x-png	.pot	application/ vnd.ms- powerpoint
.ppa	application/ vnd.ms- powerpoint	.ppm	application/x-ppm
.pps	application/ vnd.ms- powerpoint	.ppt	application/ vnd.ms- powerpoint
.ppt	application/x-ppt	.pr	application/x-pr
.prf	application/pics- rules	.prn	application/x-prn
.prt	application/x-prt	.ps	application/x-ps
.ps	application/ postscript	.ptn	application/x-ptn
.pwz	application/ vnd.ms- powerpoint	.r3t	text/vnd.rn- realtext3d
.ra	audio/vnd.rn- realaudio	.ram	audio/x-pn- realaudio
.ras	application/x-ras	.rat	application/rat- file
.rdf	text/xml	.rec	application/ vnd.rn-recording
.red	application/x-red	.rgb	application/x-rgb
.rjs	application/ vnd.rn- realsystem-rjs	.rjt	application/ vnd.rn- realsystem-rjt
.rlc	application/x-rlc	.rle	application/x-rle
.rm	application/ vnd.rn-realmedia	.rmf	application/ vnd.adobe.rmf

File Name Extension	Content-Type	File Name Extension	Content-Type
.rmi	audio/mid	.rmj	application/ vnd.rn- realsystem-rmj
.rmm	audio/x-pn- realaudio	.rmp	application/ vnd.rn- rn_music_package
.rms	application/ vnd.rn-realmedia- secure	.rmvb	application/ vnd.rn-realmedia- vbr
.rmx	application/ vnd.rn- realsystem-rmx	.rnx	application/ vnd.rn-realplayer
.rp	image/vnd.rn- realpix	.rpm	audio/x-pn- realaudio-plugin
.rsml	application/ vnd.rn-rsml	.rt	text/vnd.rn- realtext
.rtf	application/ msword	.rtf	application/x-rtf
.rv	video/vnd.rn- realvideo	.sam	application/x-sam
.sat	application/x-sat	.sdp	application/sdp
.sdw	application/x-sdw	.sit	application/x- stuffit
.slb	application/x-slb	.sld	application/x-sld
.slk	drawing/x-slk	.smi	application/smil
.smil	application/smil	.smk	application/x-smk
.snd	audio/basic	.sol	text/plain
.sor	text/plain	.spc	application/x- pkcs7-certificates
.spl	application/ futuresplash	.spp	text/xml
.ssm	application/ streamingmedia	.sst	application/ vnd.ms- pki.certstore
.stl	application/ vnd.ms-pki.stl	.stm	text/html
.sty	application/x-sty	.svg	text/svg+xml

File Name Extension	Content-Type	File Name Extension	Content-Type
.swf	application/x-shockwave-flash	.tdf	application/x-tdf
.tg4	application/x-tg4	.tga	application/x-tga
.tif	image/tiff	.tif	application/x-tif
.tiff	image/tiff	.tld	text/xml
.top	drawing/x-top	.torrent	application/x-bittorrent
.tsd	text/xml	.txt	text/plain
.uin	application/x-icq	.uls	text/iuls
.vcf	text/x-vcard	.vda	application/x-vda
.vdx	application/vnd.visio	.vml	text/xml
.vpg	application/x-vpeg005	.vsd	application/vnd.visio
.vsd	application/x-vsdx	.vss	application/vnd.visio
.vst	application/vnd.visio	.vst	application/x-vst
.vsw	application/vnd.visio	.vsx	application/vnd.visio
.vtx	application/vnd.visio	.vxml	text/xml
.wav	audio/wav	.wax	audio/x-ms-wax
.wb1	application/x-wb1	.wb2	application/x-wb2
.wb3	application/x-wb3	.wbmp	image/vnd.wap.wbmp
.wiz	application/msword	.wk3	application/x-wk3
.wk4	application/x-wk4	.wkq	application/x-wkq
.wks	application/x-wks	.wm	video/x-ms-wm
.wma	audio/x-ms-wma	.wmd	application/x-ms-wmd
.wmf	application/x-wmf	.wml	text/vnd.wap.wml
.wmv	video/x-ms-wmv	.wmx	video/x-ms-wmx

File Name Extension	Content-Type	File Name Extension	Content-Type
.wmz	application/x-ms-wmz	.wp6	application/x-wp6
.wpd	application/x-wpd	.wpg	application/x-wpg
.wpl	application/vnd.ms-wpl	.wq1	application/x-wq1
.wr1	application/x-wr1	.wri	application/x-wri
.wrk	application/x-wrk	.ws	application/x-ws
.ws2	application/x-ws	.wsc	text/scriptlet
.wsdl	text/xml	.wvx	video/x-ms-wvx
.xdp	application/vnd.adobe.xdp	.xdr	text/xml
.xfd	application/vnd.adobe.xfd	.xdf	application/vnd.adobe.xdf
.xhtml	text/html	.xls	application/vnd.ms-excel
.xls	application/x-xls	.xlw	application/x-xlw
.xml	text/xml	.xpl	audio/scpls
.xq	text/xml	.xql	text/xml
.xquery	text/xml	.xsd	text/xml
.xsl	text/xml	.xslt	text/xml
.xwd	application/x-xwd	.x_b	application/x-x_b
.sis	application/vnd.symbian.install	.sisx	application/vnd.symbian.install
.x_t	application/x-x_t	.ipa	application/vnd.iphone
.apk	application/vnd.android.package-archive	.xap	application/x-silverlight-app
.zip	application/zip	.rar	application/x-rar-compressed

## User-Defined Object Metadata

You can add the user-defined metadata whose name starts with **x-obs-meta-** for easy object management. When you retrieve or query the metadata of the object,

the added user-defined metadata will be returned in the response header. The user-defined metadata is limited to 8 KB in size. To measure the total size of user-defined metadata, sum the number of bytes in the UTF-8 encoding of each key and value.

The user-defined metadata keys are case insensitive, but are stored in lowercase by OBS. The key values are case sensitive.

The following is an example.

```
PUT /key HTTP/1.1
Host: bucket01.obs.myhuaweicloud.com
x-obs-meta-Test1: Test Meta1
```

```
HEAD /Key HTTP/1.1
Host: bucket01.obs.myhuaweicloud.com
x-obs-meta-test1: Test Meta1
```

Both user-defined metadata keys and their values must conform to US-ASCII standards. If non-ASCII or unrecognizable characters are necessary, they must be encoded or decoded in URL or Base64 on the client side, because the server side does not perform any decoding.

## How to Use

You can use OBS Console, APIs, SDKs, or obsutil to customize object metadata.

Tool	Reference
OBS Console	<a href="#">Configuring Object Metadata</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Modifying Object Metadata</a>
obsutil	<a href="#">Uploading an Object</a>

## 4.5.6 Restoring Objects from the Archive Storage

### Scenarios

You must restore an object in the Archive storage class before you can download it, configure its metadata, or access it with a URL.

Archive objects can be restored at an expedited speed within 1 to 5 minutes or at a standard speed within 3 to 5 hours.

### Constraints

- If an Archive object is being restored, you cannot suspend or delete the restore task.
- You cannot re-restore an object that is in the **Restoring** state.

- After an object is restored, an object copy in the Standard storage class will be generated. This way, there is an Archive object and a Standard object copy in the bucket at the same time. After an Archive object is restored, the object status displays **Restored**, and the generated object copy in Standard storage class is not displayed in the bucket.

During the restore validity period, you will be billed for the space taken up by both the object and its copy. The copy will be automatically deleted upon expiration of its validity period.

## How to Use

You can use OBS Console, APIs, SDKs, OBS Browser+, or obsutil to restore Archive objects.

Tool	Reference
OBS Console	<a href="#">Restoring Objects from the Archive Storage</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Restoring Archive Objects</a>
OBS Browser+	<a href="#">Restoring a File</a>
obsutil	<a href="#">Restoring Objects from OBS Archive</a>

## 4.5.7 Direct Reading

### Scenarios

You can enable direct reading if you want to obtain Archive objects immediately or if there are interface compatibility issues when OBS is interconnected with other systems.

With direct reading enabled for a bucket, you can download objects in the Archive storage class without restoring them first.

You can enable direct reading when creating a bucket or enable it for an existing bucket.

### How to Use

You can use OBS Console or APIs to configure direct reading.

Tool	Reference
OBS Console	<a href="#">Direct Reading</a>
API	<a href="#">Setting the Direct Reading Policy for Archive Objects in a Bucket</a>

## 4.6 Deleting Objects

### 4.6.1 Deleting an Object

#### Scenarios

You can delete unwanted files or folders to save storage space and reduce costs.

---

**NOTICE**

If versioning is not enabled for a bucket, deleted objects cannot be recovered. Exercise caution when performing this operation.

---

#### Manually or Automatically Deleting Objects

You can manually delete objects anytime. Alternatively, you can use [lifecycle rules](#) to periodically delete objects from a bucket or clear all objects in a bucket at a time.

In big data scenarios, parallel file systems usually have deep directory levels and each directory has a large number of files. In such case, deleting directories from parallel file systems may fail due to timeout. To address this problem, you are advised to delete directories in either of the following ways:

1. On the Hadoop client that has OBSA, an OBS client plugin, embedded, run the `hadoop fs -rmr obs://{Name of a parallel file system}/{Directory name}` command.
2. Configure [a lifecycle rule](#) for directories so that they can be deleted in background based on the preset lifecycle rule.

#### Deleting an Object or a Batch of Objects

You can delete one or more objects at a time. For details about how to delete objects in batches, see [Does OBS Support Batch Deletion of Objects?](#)

** NOTE**

The batch deletion performance is negatively correlated with the number of objects in a single request. When it comes to QPS, deleting  $N$  objects is counted as  $N$  operations. If a large number of objects named with prefixes in lexicographic order are deleted, lots of requests may be concentrated in a specific partition, which results in hot access. This limits the request rate in the hot partition and increases access delay.

To address this problem, you can reduce the number of objects in a single batch deletion request, initiate more concurrent requests, and name objects with random prefixes.

#### Object Deletion Mechanism When Versioning Is Enabled

When versioning is enabled, OBS uses different deletion methods for different objects.

- Deleting a file or folder does not delete it permanently. The deleted file or folder will be retained in the **Deleted Objects** list and marked with the **Delete Marker**.
  - If you want to delete the file or folder permanently, you need to delete it from the **Deleted Objects** list.
  - To recover a deleted file, you can cancel the deletion by the **Undelete** operation. For details, see [Undeleting an Object](#).
- Deleting a version of an object will permanently delete that version. If the deleted version is the latest one, the next latest version becomes the latest version.

When versioning is enabled, files in the **Deleted Objects** list also have multiple versions. Note the following points when deleting different versions of files:

- If you delete a version with the **Delete Marker**, it actually recovers that specific version instead of permanently deleting it. For details, see [Undeleting an Object](#).
- If you delete a version without the **Delete Marker**, that specific version is deleted permanently. Even if the object is recovered later, this version will not be recovered.

For more information, see [Object Recovery Mechanism When Versioning Is Enabled](#).

## How to Use

You can use OBS Console, APIs, SDKs, OBS Browser+, or obsutil to delete objects.

Tool	Reference
OBS Console	<a href="#">Deleting a File or Folder</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Deleting an Object</a>
OBS Browser+	<a href="#">Deleting a File or Folder</a>
obsutil	<a href="#">Deleting an Object</a>

## 4.6.2 Undeleting an Object

### Scenarios

If versioning is not enabled for a bucket, deleted objects cannot be recovered. If versioning is enabled, you can recover deleted objects through the **Undelete** operation. For details about versioning, see [Versioning](#).



## Object Recovery Mechanism When Versioning Is Enabled

When a bucket has the versioning function enabled, deleting a file from the **Objects** list does not permanently delete it. The deleted file will be retained with the **Delete Marker** in the **Deleted Objects** list. You can recover a deleted object by the **Undelete** operation.

When performing the **Undelete** operation, note the following points:

1. You can only undelete deleted files but not folders.  
After you undelete a deleted file, the file is recovered and will appear in the **Objects** list. Then you can perform basic operations on the file as you normally do on other objects. If the file was stored in a folder before the deletion, it will be recovered to its original path after you undelete it.
2. Deleted files in the **Deleted Objects** also have multiple versions. When deleting different versions of files, note the following points:
  - If you delete a version with the **Delete Marker**, it actually recovers that specific version instead of permanently deleting it.
  - If you delete a version without the **Delete Marker**, that specific version is deleted permanently. Even if the object is recovered later, this version will not be recovered.
3. At least one version without the **Delete Marker** exists in the **Deleted Objects** list. Otherwise, the deletion cannot be canceled.

### How to Use

You can use OBS Console to undelete objects.

Tool	Reference
OBS Console	<a href="#">Undeleting a File</a>

## 4.6.3 Managing Fragments

### Scenarios

Fragments refer to incomplete data in a bucket. Data is uploaded to OBS in multipart mode. In the following common and other scenarios, an upload fails and fragments are generated. You need to clear these fragments to free up storage space.

- The network is in poor conditions, and the connection to the OBS server is interrupted frequently.
- The upload task is manually suspended.
- The device is faulty.
- The device is powered off suddenly.

Fragments in OBS consume storage space and are billed based on price rates of storage space. If fragments are generated due to interruptions or failures of

multipart upload tasks, you can resume such tasks so that fragments will be cleared, or you can directly clear the fragments to save storage space.

 **NOTE**

- If a bucket fails to be deleted, check whether all fragments have been deleted. If not, delete all fragments from the bucket.
- If no object exists in the bucket but the fee is still being deducted, check whether there are fragments in the bucket. If yes, delete the fragments to reduce storage costs.

## How to Use

**Table 4-4** Methods for deleting fragments

Tool	Method
OBS Console	OBS Console allows you to batch delete all selected fragments with one click. For details, see <a href="#">Managing Fragments</a> .
OBS Browser+	You can delete all of the selected fragments in a bucket in a batch. Alternatively, you can click <b>Delete All</b> above the list to delete all fragments.
obsutil	With obsutil, you can delete the failed or interrupted multipart upload task to delete fragments generated by the task. If a bucket has more than one multipart upload tasks, you need to delete all the multipart upload tasks to delete all fragments in the bucket. For details, see <a href="#">Deleting a Multipart Upload Task</a> .
API	You can delete fragments from a bucket through the following procedure: <ol style="list-style-type: none"><li>1. Use the <a href="#">Listing Initiated Multipart Uploads in a Bucket</a> operation to list all the multipart upload tasks and obtain their upload IDs.</li><li>2. Use the <a href="#">Canceling a Multipart Upload Task</a> operation to cancel the multipart upload tasks. After these tasks are canceled, all fragments generated by them can be deleted.</li></ol>
SDKs	With OBS SDK, fragments are generated when parts of a multipart task are not merged. You can cancel the task to delete generated fragments. The procedure is as follows: <ol style="list-style-type: none"><li>1. Use the <b>ObsClient.listMultipartUploads</b> interface to list all multipart upload tasks and obtain their UploadId.</li><li>2. Use the <b>ObsClient.abortMultipartUpload</b> interface to cancel the multipart upload task so that the generated fragments will be cleared.</li></ol>

# 5 Permissions Control

## Scenarios

By default, OBS resources (buckets and objects) are private. Only resource owners can access their OBS resources. Without authorization, other users cannot access your OBS resources. OBS permission control means to grant permissions to other accounts or IAM users by editing access policies. For example, if you have a bucket, you can authorize another IAM user to upload objects to your bucket. You can also open buckets to the public, so that anyone can access your buckets over the Internet. OBS offers multiple methods to help you to assign resource permissions to others. Resource owners can formulate different permissions control policies based on service requirements to ensure data security.

## How to Use

You can use OBS Console, APIs, SDKs, OBS Browser+, or obsutil to configure permissions.

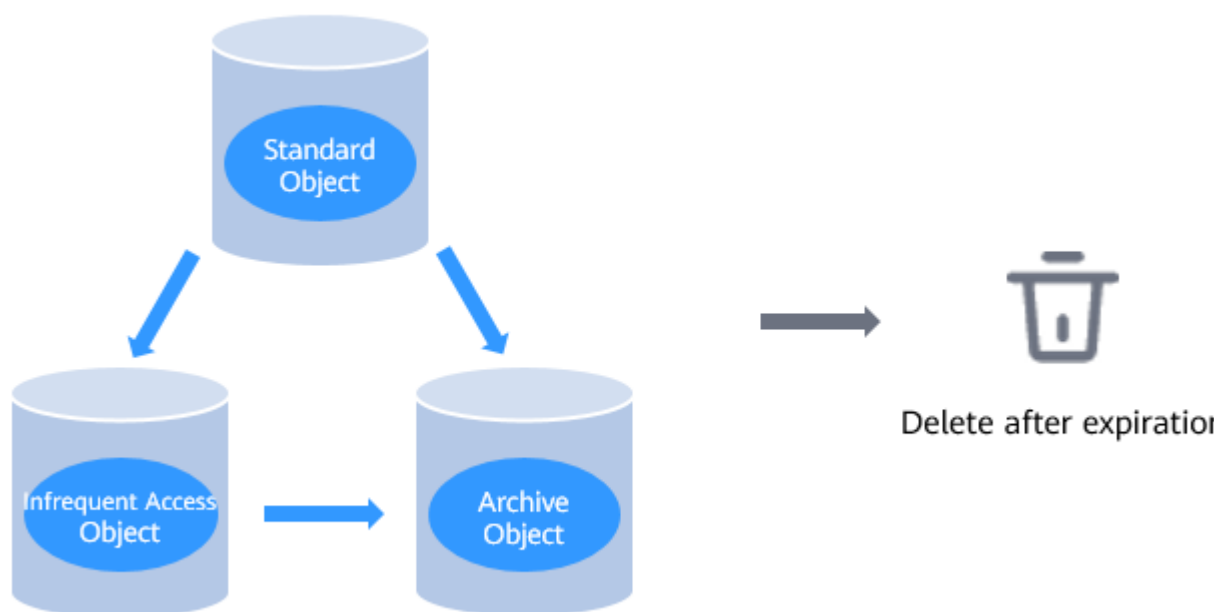
Tool	Reference
OBS Console	<a href="#">Permission Control</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Permissions and Supported Actions</a> <a href="#">Configuring a Bucket Policy</a> <a href="#">Configuring a Bucket ACL</a> <a href="#">Configuring an Object ACL</a>
OBS Browser+	-
obsutil	<a href="#">Setting Bucket Properties</a> <a href="#">Setting Object Properties</a>

# 6 Data Management

## 6.1 Lifecycle Management

Lifecycle management for OBS includes periodically deleting objects from buckets and transitioning between object storage classes based on pre-defined rules.

**Figure 6-1** Lifecycle management diagram



### Scenarios

You may configure lifecycle rules to:

- Periodically delete logs that are only meant to be retained for a specific period of time (a week or a month).
- Transition documents that are seldom accessed to the Infrequent Access or Archive storage class or delete them.

- Store some types of data in OBS for archive purposes, such as digital media, financial and medical records, raw genome sequence data, long-term database backups, and data that must be retained for regulatory compliance.
- Schedule the deletion of a large number of files from a bucket. Manually deleting objects is time-consuming, and only a limited number of objects can be deleted at a time.

You can define lifecycle rules to identify objects in the scenarios above and further manage their lifecycles.

Lifecycle configuration can be added to a bucket with versioning enabled or disabled. By default, versioning is disabled. You can enable it. A version-enabled bucket keeps a current object version and zero or more non-current object versions. You can use versioning and lifecycle configurations together to help reduce storage costs. The predefined lifecycle configuration actions can facilitate management over the lifecycle of the current object versions and non-current object versions.

## Precautions

- There is no limit on the number of lifecycle rules in a bucket, but the total size of XML descriptions about all lifecycle rules in a bucket cannot exceed 20 KB.
- A maximum of 20 lifecycle rules can be configured for a parallel file system.
- A lifecycle rule name contains only uppercase or lowercase letters, digits, periods (.), underscores (\_), and hyphens (-).
- The minimum storage duration is 90 days for Archive storage, and 30 days for Infrequent Access storage. After an object is transitioned to the Archive storage class, if it stays in this storage class for less than 90 days, you still need to pay for a full 90 days.
- Restrictions on storage class transition:
  - Only transitions from the Standard storage class to the Infrequent Access storage class are supported. To transition objects from Infrequent Access to Standard, you must manually operate it.
  - Only transitions from the Standard or Infrequent Access storage class to the Archive storage class are supported. To transition objects from Archive to Standard or Infrequent Access, you must restore the archived objects first and then manually transition their storage class.
- Object deletion upon expiration:

You can use the **Expiration** action a lifecycle rule to expire an object when the object reaches the end of its lifetime. In a versioning-enabled bucket, you can also use the **NoncurrentVersionExpiration** action to expire non-current object versions.

  - In a bucket with versioning disabled, the **Expiration** action permanently deletes objects.
  - In a bucket with versioning enabled (or suspended), the **Expiration** action retains the current version as a non-current version by adding a delete marker, which then becomes the current version. The **NoncurrentVersionExpiration** action permanently deletes non-current object versions.

- Beside buckets, you can also configure the expiration time for objects during object uploads. The expiration time of objects prevails against that of buckets. After an object expires, OBS adds the object to the deletion queue and delete it asynchronously. This may cause the deletion time to be later than the expiration time. After an object expires, you will not be charged for the related storage duration fee.  
To query the planned expiration time of an object, you can call the GET object API or HEAD object API. These APIs return response headers that provide expiration information about the object.
- After a lifecycle rule is modified, the modification does not apply to the objects that have already met the configured conditions. For example, the original lifecycle rule causes objects meeting the configured conditions expire and then be deleted after they are stored for one day. Under this rule, the objects uploaded on January 1, 2021 will be deleted on January 3, 2021. On January 3, 2021, if you change the lifecycle rule to delete objects after they are stored for seven days, the objects uploaded on January 1, 2021 will still be deleted on January 3, 2021, but this modification will apply to those uploaded on and after January 2, 2021.

## Lifecycle Rules

Lifecycle rules have the following key elements:

- Policy: Configure a lifecycle rule that takes effect on specified objects.
  - By prefix: You can specify an object name prefix, so the lifecycle rule will take effect on objects share the same prefix.
  - Entire bucket: You can specify an entire bucket, so the lifecycle rule will take effect on all objects in the bucket.
- Time: a scheduled time when object storage class is transitioned  
You can specify the number of days after which objects that have been last updated and meet specified conditions are automatically transitioned to Infrequent Access or Archive, or are automatically deleted upon expiration.
  - Transition to Infrequent Access: You can specify the number of days after which objects that have been last updated and meet specified conditions are automatically transitioned to Infrequent Access.
  - Transition to Archive: You can specify the number of days after which objects that have been last updated and meet specified conditions are automatically transitioned to Archive.
  - Deleted upon expiration: You can specify the number of days after which objects that have been last updated and meet specified conditions are automatically deleted.

The lifecycle rule time has the following restrictions:

- If only one transition is involved, the time should be at least one day later.
- If multiple transitions are involved, the latter transition should be at least one day longer than the former transition.

For example, if you want a lifecycle rule to change object storage class and delete expired objects, you can set the transition time to 23 days later and the deletion time to at least 24 days later.

 NOTE

The minimum storage duration is 90 days for Archive storage, and 30 days for Infrequent Access storage. If objects are deleted by a lifecycle rule before they have been stored for this minimum duration, you still need to pay for a full 30 or 90 days.

Assume that you have an object in the Standard storage class stored in a bucket and you no longer plan to perform any operations on the object. The lifecycle rule applied to this object changes the object from Standard to Infrequent Access three days later, then to Archive on the fourth day, and finally deletes the object when it expires seven days later. In this case, you will be charged for three days of Standard storage, 30 days of Infrequent Access storage, and 90 days of Archive storage.

## Lifecycle Configuration Elements

You can set the lifecycle configuration format to XML. The configuration contains one or more lifecycle rules.

Each rule consists of the following contents:

- Metadata, specifying the rule ID and whether the rule is enabled or disabled. If the rule is disabled, OBS does not perform any actions defined in the rule.
- Filtering criteria, identifying objects on which lifecycle rules are applied. You can set the object name prefix to be the filtering criteria.
- When (a date or a time period) will one transition or expiration action be performed on objects in the lifecycle.

Configuration examples:

### Example 1: Lifecycle configuration for a bucket with versioning disabled

By default, versioning is disabled for buckets. In such case, each object in the bucket has only one version.

Assume that versioning is disabled for your bucket. If you want Standard objects whose name starts with **documents/** to transition to the Infrequent Access storage class 30 days after they are uploaded, then transition from Infrequent Access to Archive 60 days after they are uploaded, and finally to be deleted one year after they are uploaded, you can add the following lifecycle configuration for your bucket. This configuration includes the **Transition** and **Expiration** actions and applies to objects whose key prefix is **documents** (specified in the **Prefix** element).

```
<LifecycleConfiguration>
  <Rule>
    <ID>sample-rule</ID>
    <Prefix>documents/</Prefix>
    <Status>Enabled</Status>
    <Expiration>
      <Days>365</Days>
    </Expiration>
    <Transition>
      <Days>30</Days>
      <StorageClass>WARM</StorageClass>
    </Transition>
    <Transition>
      <Days>60</Days>
      <StorageClass>COLD</StorageClass>
    </Transition>
  </Rule>
</LifecycleConfiguration>
```

### Example 2: Lifecycle configuration for a bucket with versioning enabled

You can enable versioning for buckets. If versioning is enabled for a bucket, the bucket will retain the current object version and its non-current object versions. For details, see [Versioning](#). Versioning control enables you to maintain the history records of objects and lifecycle management allows you to control the retention of object versions as well as the storage class transition.

For a bucket with versioning enabled, lifecycle configuration provides multiple predefined actions that can be used to manage non-current object versions. In this example, the lifecycle configuration has a rule that specifies two operations (**NoncurrentVersionTransition** and **NoncurrentVersionExpiration**) for objects whose key prefix is **logs/**. The **NoncurrentVersionTransition** action transitions objects to Infrequent Access and then to Archive 30 days and 60 days respectively after they become non-current versions. The **NoncurrentVersionExpiration** action permanently deletes objects 180 days after they become non-current versions.

```
<LifecycleConfiguration>
  <Rule>
    <ID>sample-rule</ID>
    <Prefix>logs/</Prefix>
    <Status>Enabled</Status>
    <NoncurrentVersionExpiration>
      <NoncurrentDays>180</NoncurrentDays>
    </NoncurrentVersionExpiration>
    <NoncurrentVersionTransition>
      <NoncurrentDays>30</NoncurrentDays>
      <StorageClass>WARM</StorageClass>
    </NoncurrentVersionTransition>
    <NoncurrentVersionTransition>
      <NoncurrentDays>60</NoncurrentDays>
      <StorageClass>COLD</StorageClass>
    </NoncurrentVersionTransition>
  </Rule>
</LifecycleConfiguration>
```

You can use the predefined **NoncurrentVersionTransition** and **NoncurrentVersionExpiration** actions to manage non-current versions in your bucket.

Generally, each lifecycle rule consists of the following parts:

- Metadata: specifying the rule ID (**<ID>** element) and whether the rule is enabled or disabled (**<Status>** element). If the rule is disabled, OBS does not perform any actions defined in the rule.
- Prefix (**<Prefix>** element), which identifies objects to which the rule applies.
- Actions that you want to perform on the specified objects (such as **NoncurrentVersionTransition** and **NoncurrentVersionExpiration** mentioned in the previous example). For each action, you must specify a date or time period.

Elements in the lifecycle configuration rule are described as follows:

- **Element ID**  
A lifecycle configuration can have up to 1,000 rules. The ID element uniquely identifies a rule.
- **Element Status**  
The value can be **Enabled** or **Disabled**. If a rule is disabled, OBS does not perform any actions defined in the rule.



- Element **Prefix**

A lifecycle rule applies to one or more objects with the name prefix specified in the rule. Suppose you have the following objects:

- logs/day1
- logs/day2
- logs/day3
- ExampleObject.jpg

If you set **Prefix** to **ExampleObject.jpg**, the rule applies only to object **ExampleObject.jpg**. If you set **Prefix** to **logs/**, the rule applies to the objects whose name starting with **logs/**. If you leave **Prefix** blank, the rule applies to all objects in your bucket.

- Element **Action**

You can specify predefined actions in the lifecycle rule to perform them on your buckets in the lifecycle. The predefined actions include **Transition**, **Expiration**, **NoncurrentVersionTransition**, and **NoncurrentVersionExpiration**. The effect of these operations depends on the versioning status of your bucket.

 **NOTE**

By default, versioning is disabled for buckets. You can enable it for your bucket as needed, so that each object in this bucket can have a current version and zero or more non-current versions. You can also suspend versioning. For details, see [Versioning](#).

- Action **Transition**

This action archives objects by changing object storage class to Infrequent Access or Archive. When the date or time period specified in the lifecycle of an object is reached, OBS transitions the object as configured.

- For buckets with versioning disabled, the **Transition** action changes the object storage class to Infrequent Access or Archive.
- For buckets with versioning enabled and suspended, the **Transition** action changes the storage class of the current object version to Infrequent Access or Archive. This action does not affect the non-current versions of the object.

- Action **Expiration**

This action expires objects identified in the rule. Objects become unavailable once they expire. Whether the expired objects will be permanently deleted depends on the versioning status of the bucket.

 **NOTE**

The **Expiration** action does not delete incomplete multipart uploads.

- For buckets with versioning disabled, the **Expiration** action deletes objects permanently and the deleted objects cannot be restored.
- For buckets with versioning enabled, this action applies only to current object versions, instead of non-current object versions. This action does not permanently delete current object versions. It retains the current version as a non-current version by adding a delete marker to it. This action will not be performed on current object versions that already have the delete markers. If the current object version is the only version of the object and has a delete marker, OBS will delete the current object

version. Clearing a delete marker may take a while, because OBS needs to confirm that the delete marker is the only object version.

If you initiate a GET request on an object whose current version has the delete marker without specifying the version ID, OBS will identify the object as a deleted one and return error **404 Object Not Found**. But you can specify the version ID in the GET request to recover the deleted object.

For example, you can set a rule to make the object named **photo.gif** expire 5 days after it is uploaded. If **photo.gif** is created at 10:30 UTC on January 1, 2016, the expiration rule will be executed at a time point after 00:00 UTC (five days after object creation) on January 7, 2016. The time will not be later than 23:59 UTC on January 7, 2016. For a bucket with versioning disabled, a deletion operation permanently deletes **photo.gif**. For a bucket with versioning enabled, after the expiration rule is executed, **photo.gif** (version **111111**) is still stored in the bucket and can be accessed if needed, but the current version (version **4857693**) of the object has a delete marker. The original object **photo.gif** turns to be a non-current version. For details about how a delete marker works, see [Versioning](#).

For buckets with versioning suspended, OBS will create delete markers for expired objects whose version ID is **null**. Any existing **null** versions will be overwritten by new **null** versions, and data associated with this version cannot be restored.

#### **Actions specific to buckets with versioning control enabled (or suspended)**

The **Transition** and **Expiration** lifecycle actions can manage the lifecycle of current object versions. The **NoncurrentVersionTransition** and **NoncurrentVersionExpiration** actions can manage the lifecycle of noncurrent object versions.

The following lifecycle configuration actions can be performed only on buckets with versioning control enabled (or suspended). In a bucket with versioning control enabled, an object may have multiple versions, including a current version and zero or more non-current versions. You can use these actions to request OBS to perform specific operations on non-current object versions. These actions do not affect current object versions.

**NoncurrentVersionTransition**: Specifies the time period after which non-current versions will be transitioned from Standard to Infrequent Access or Archive.

**NoncurrentVersionExpiration**: Specifies the time period after which non-current object versions will be permanently deleted. The deleted objects cannot be recovered.

For example, if you want to enable a five-day period to correct any accidental deletion or overwriting, you can configure an expiration rule so that the object can be deleted 5 days after it has become a non-current version.

At 10:30 (UTC time) on January 1, 2016, you created object **photo.gif** with version ID **111111**. At 10:30 (UTC time) on January 2, 2016, you accidentally deleted this object and OBS created a delete marker with version ID **4857693**. In the next five days, you were allowed to recover the original object **photo.gif** with version ID **111111**.

At 00:00 (UTC time) on January 8, 2016, 5 days after object **photo.gif** with version ID **111111** became non-current, the **NoncurrentVersionExpiration** action permanently deleted this object version.

### How does OBS calculate how long an object has become non-current?

In a bucket with versioning enabled, an object can have multiple versions. There is always one current version and zero or more non-current versions. Each time an object is uploaded, the current version is saved as a non-current version, and the newly-uploaded version (the successor) becomes the current version. To determine the number of days an object version has become non-current, OBS checks when its successor was created. OBS uses the number of days since its successor was created as the number of days an object has been non-current.

### Restoring a non-current version using the lifecycle configuration

You can use either of the following methods to restore an object to a non-current version:

1. Copy a non-current object version of the object to the same bucket. The copied version will become the current version, and all object versions are reserved.
2. Permanently delete the current version of the object. Deleting the current version of an object actually turns a non-current version into the current version.

You are advised to use the first method. Due to the consistency syntax of OBS, before the communication mode is changed, a current version that is permanently deleted may not disappear (OBS may be unaware of this deletion action). Meanwhile, the expiration action you configured for non-current versions may permanently delete non-current object versions, including those you want to restore. Therefore, copying a non-current version, as described in the first method, is more secure.

**Table 6-1** lists the relationship between the lifecycle configuration rule and versioning status of a bucket.

**Table 6-1** Lifecycle actions and bucket versioning status

Action	Bucket with Versioning Not Enabled	Bucket with Versioning Enabled	Bucket with Versioning Suspended
Transition (Performed when a date or time period specified in the lifecycle rule is reached.)	Objects can be transitioned to Infrequent Access or Archive.	If the version is the current version and is not a delete marker, it can be transitioned to Infrequent Access or Archive.	Same action as a versioning-enabled bucket.

Action	Bucket with Versioning Not Enabled	Bucket with Versioning Enabled	Bucket with Versioning Suspended
Expiration (Performed when a date or time period specified in the lifecycle rule is reached.)	The <b>Expiration</b> action deletes the object, and the deleted object cannot be restored.	If the version is not a delete marker, this action inserts a delete marker that becomes the current version. The existing current version turns into a non-current version.	For buckets with versioning suspended, OBS will create delete markers for expired objects whose version ID is <b>null</b> . Any existing <b>null</b> versions will be overwritten by new <b>null</b> versions, and data associated with this version cannot be restored.
NoncurrentVersionTransition (Performed when the object has become non-current for the specified number of days.)	The NoncurrentVersionTransition action has no effect.	If the version is not a delete marker or the current version, it can be transitioned to Infrequent Access or Archive.	Same action as a versioning-enabled bucket.
NoncurrentVersionExpiration (Performed when the object has become non-current for the specified number of days.)	The NoncurrentVersionExpiration action has no effect.	The NoncurrentVersionExpiration action deletes the non-current version of the object, and the deleted object cannot be recovered.	Same action as a versioning-enabled bucket.

### Date-based lifecycle rules

You can specify the execution dates for **Transition** and **Expiration** actions. The dates must conform to the ISO 8601 standards and the exact time is always 00:00 (UTC time). If you specify the time to a past date, all applicable objects will be executed immediately.

You cannot create date-based lifecycle rule on OBS Console.

---

**WARNING**

A lifecycle action with a date specified is not a one-off action. Even if the date has passed, OBS will adopt this action as long as the lifecycle is enabled.

Assume you have specified a date for performing the **Expiration** action to delete all objects (without setting any filtering criteria). On the specified date, OBS makes all objects in the bucket expire. In addition, OBS continues to make all new objects created in the bucket expire. To terminate the **Expiration** action, you have to delete this action from your lifecycle configuration, disable the rule, or delete the rule from the lifecycle configuration.

---

**Time period-based lifecycle rule**

You can specify how many days after an object is created will the **Transition** or **Expiration** action will be performed on the object. After the number of days is specified, OBS starts to calculate the time from 00:00 (UTC time) on the next day. For example, you created an object at 2016-01-15 10:30 (UTC time) and you specified that objects would be changed 3 days after creation, the object would be transitioned at 2016-01-19 00:00 (UTC time).

---

**NOTICE**

OBS only records the last modification date for each object. On OBS Console, you can view the last modification time (**LastModified**) of an object on the object properties page. After an object is created, the date is the creation date. If the object is replaced, the date will also change.

---

When using the **NoncurrentVersionTransition** or **NoncurrentVersionExpiration** action, you can specify after how many days since an object changes to a non-current version (due to overwrite or deletion) will the action be performed on the object.

When you use the **NoncurrentVersionTransition** and **NoncurrentVersionExpiration** operations to specify the number of days, OBS adds the number of days specified in the rule to the time when the object version becomes a non-current version. Then the operation time is calculated starting from 00:00 (UTC time) of the next day. For example, the current version of an object in a bucket is created at 10:30 UTC on January 1, 2016. If the object version becomes a non-current version at 10:30 UTC on January 15, 2016 and you specify three days in the conversion rule, then the date for changing object storage class is calculated as 00:00 UTC on January 19, 2016.

**NOTE**

When configuring the lifecycle rules, within a rule and for rules whose prefixes have inclusion relationship, **Date** or **Days** of **Transition** and **Expiration** must be consistent.

## How to Use

You can use OBS Console, APIs, or SDKs to configure lifecycle rules.

Tool	Reference
OBS Console	<a href="#">Configuring a Lifecycle Rule</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Configuring Bucket Lifecycle Rules</a> <a href="#">Obtaining Bucket Lifecycle Configuration</a> <a href="#">Deleting Lifecycle Rules</a>

## 6.2 Tags

### Scenarios

Tags help you to identify your cloud resources. When you have many cloud resources of the same type, you can use tags to classify them by dimension (for example, by purpose, owner, or environment) for usage or cost analysis.

In OBS, you can use tags to identify and classify OBS buckets.

If you add tags to a bucket, service detail records (SDRs) generated for it will be labeled with these tags. You can classify SDRs by tag for cost analysis. For example, if you have an application that uploads its running data to a bucket, you can tag the bucket with the application name. In this manner, the costs can be analyzed using tags in SDRs.

### Constraints

- A tag is structured as a key-value pair that is case sensitive.
- A bucket can have a maximum of 10 tags. Each tag has only one key and one value.
- Each key must be unique among all tags of a bucket, whereas the tag values can be repetitive or left blank.
- It takes approximately three minutes for an added tag to take effect.

### How to Use

You can configure tags using OBS Console.

Tool	Reference
OBS Console	<a href="#">Configuring Tags for a Bucket</a>

## 6.3 Bucket Inventory

### Scenarios

A bucket inventory can list objects in a bucket, save the related object information in CSV files, and deliver the CSV files to the bucket specified for storing bucket inventory files. In this manner, you can easily manage objects in a bucket. A source bucket can also be the destination bucket.

- A bucket inventory file can contain the following object related information: versions, sizes, storage classes, tags, encryption statuses, and last modification time.
- You can encrypt bucket inventory files in the SSE-KMS mode.
- You can set the frequency (daily or weekly) for generating bucket inventory files.
- You can also specify a bucket to store the generated bucket inventory files.

### Constraints

#### Bucket versions

- Inventories can be generated only for OBS 3.0 buckets, but they can be stored in either OBS 3.0 or OBS 2.0 buckets.

#### Number of bucket inventories

- A bucket can have a maximum of 10 inventories.

#### Source and destination buckets

- The source bucket (for which a bucket inventory rule is configured) and the destination bucket (where the generated inventory files are stored) must belong to the same account.
- The source bucket and the destination bucket must be in the same region.
- Default encryption cannot be enabled for the destination bucket configured for storing inventory files.

#### Functions

- Inventory files must be in the CSV format.
- OBS can generate inventory files for all objects in a bucket or a group of objects whose names begin with the same prefix.
- If a bucket has multiple inventory rules, these rules must not overlap.
  - If a bucket already has an inventory rule for the entire bucket, new inventory rules that filter objects by prefixes cannot be created. If you need an inventory rule that covers only a subset of objects in the bucket, first delete the inventory rule configured for the entire bucket.
  - If an inventory rule that filters objects by a specified prefix already exists, you cannot create an inventory rule for the entire bucket. To create an inventory rule for the entire bucket, make sure that the bucket has no other inventory rules that filter objects by specified prefixes.

- If a bucket already has an inventory rule that filters objects by the object name prefix **ab**, the filter of a new inventory rule cannot start with **a** or **abc**. Or, you can delete the existing inventory rule and create a new one that filters objects according to your needs.
- Only SSE-KMS can be used to encrypt bucket inventories.

#### Permissions

- Inventory files are delivered to the destination bucket by an OBS system user. Therefore, you need to authorize the system user the permission to write the destination bucket.

#### Others

- The bucket inventory function is offered for free, but inventory files are billed for the storage space they occupy.

## How Is a Bucket Inventory Configured?

Before the configuration, you need to briefly understand what a source bucket or a destination bucket is.

- Source bucket: A source bucket is the bucket for which an inventory is configured. The inventory lists objects stored in the source bucket.
- Destination bucket: A destination bucket is where generated inventory files are stored. A source bucket can also be the destination bucket. You can specify a name prefix for an inventory. Then generated inventory files will be named with the prefix and saved in the directory with the prefix. If you do not specify any name prefix for the inventory, the generated inventory files are stored in the root directory of the bucket.
  - Restrictions on the destination bucket
    - The destination bucket and source bucket must belong to the same tenant.
    - The destination bucket and source bucket must be in the same region.
    - A bucket policy must be configured to grant OBS the permission to write objects to the destination bucket. For details, see [Add a bucket policy for the destination bucket](#).
  - The destination bucket contains the following files:
    - A list of inventory files
    - The **Manifest** file, which contains the list of all inventory files under a certain inventory configuration. For details about the **Manifest** file, see [Manifest File](#).

#### Configuring a Bucket Inventory

You use OBS Console or call the API to configure a bucket inventory. If you configure a bucket inventory on OBS Console, a bucket policy with the required permission configuration is automatically generated for the destination bucket. If you call the API to configure the bucket inventory, you need to manually configure the bucket policy for the destination bucket.



1. Add a bucket policy for the destination bucket.

A bucket policy must be configured for the destination bucket, to grant the OBS system users the permission to write objects to the destination bucket. The format of the bucket policy is as follows. Replace **destbucket** with the actual name of the destination bucket.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Sid": "1",
      "Principal": {"Service": "obs"},
      "Resource": ["destbucket/*"],
      "Action": ["PutObject"]
    }
  ]
}
```

2. Configure a bucket inventory.

We provide multiple tools to configure a bucket inventory. For details, see [How to Use](#).

## Content in an Inventory File

The content in an inventory file can be configured when creating the inventory. For details about all possible fields, see [Table 6-2](#).

**Table 6-2** Object metadata listed in an inventory file

Metadata	Description
Bucket	Name of the source bucket
Key	The name of an object. Each object in a bucket has a unique key. (Object names in the inventory file are URL-encoded using UTF-8 character set and can be used only after being decoded.)
VersionId	Version ID of an object. If the value of <b>IncludedObjectVersions</b> in the inventory configuration is <b>Current</b> , this field is not included in the inventory file.
IsLatest	If the object version is the latest, this parameter is <b>True</b> . (If the value of <b>IncludedObjectVersions</b> in the inventory configuration is <b>Current</b> , this field is not included in the inventory file.)
IsDeleteMarker	When versioning is enabled for the source bucket, if an object is deleted, a new object metadata is generated for the object, and the <b>IsDeleteMarker</b> of the metadata is set to <b>true</b> . (If the value of <b>IncludedObjectVersions</b> in the inventory configuration is <b>Current</b> , this field is not included in the inventory file.)
Size	Object size, in bytes.
LastModifiedDate	Object creation date or last modification date

Metadata	Description
ETag	Hexadecimal digest of the object MD5. ETag is the unique identifier of the object content. It can be used to identify whether the object content is changed. For example, if ETag value is <b>A</b> when an object is uploaded and the ETag value has changed to <b>B</b> when the object is downloaded, it indicates that the object content is changed.
StorageClass	Storage class of an object
IsMultipartUploa- ded	Indicates whether an object is uploaded in the multipart mode.
ReplicationStatus	Cross-region replication status of an object
EncryptionStatus	Encryption status of an object

## Inventory File Name

The name of an inventory file is in the following format:

```
destinationPrefix/sourceBucketName/inventoryId/yyyy-MM-dd'T'HH-mm'Z'/files/UUID_index.csv
```

- **destinationPrefix**: The inventory file name prefix configured when creating the inventory rule. Inventory files generated under the rule are named after the prefix, which can facilitate the classification of inventory files. If no prefix is specified, the default prefix is **BucketInventory**.
- **sourceBucketName**: Name of the source bucket for which an inventory is configured. This field can be used to differentiate inventory files of different source buckets, if those inventory files are saved in the same destination bucket.
- **inventoryId**: If a source bucket has multiple inventory rules whose inventory files are saved in the same destination bucket, this field can be used to identify different inventory rules.
- **yyyy-MM-dd'T'HH-mm'Z'**: Start time and date for scanning the destination bucket when an inventory file is generated. Objects uploaded to the source bucket after this time may not be listed in the inventory file.
- **UUID\_index.csv**: one of the inventory files

## Manifest File

If there are a large number of objects in a bucket, multiple inventory files may be generated for a single inventory configuration. It takes some time to generate these files. For example, if there are 200,000 objects in a bucket, it takes about 1.5 minutes to generate all inventory files. One or two hours after all inventory files are generated, a **manifest.json** file will be generated. The **manifest.json** file contains information about all inventory files generated this time, including:

- **sourceBucket**: name of the source bucket
- **destinationBucket**: name of the destination bucket

- **version**: version of the inventory
- **fileFormat**: format of inventory files
- **fileSchema**: object metadata fields contained in the inventory files
- **files**: list of all inventory files
- **key**: inventory file name
- **size**: size of an inventory file, in bytes
- **inventoriedRecord**: number of records contained in an inventory file

The following is an example of a simple **manifest.json** file.

```
{
  "sourceBucket": "user001",
  "destinationBucket": "bucket001",
  "version": "2019-01-03",
  "fileFormat": "CSV",

  "fileSchema": "Bucket,Key,Size,LastModifiedDate,ETag,StorageClass,IsMultipartUploaded,ReplicationStatus,EncryptionStatus",
  "files": [
    {
      "key": "inventory/user001/test_id/2019-01-03T12-28Z/files/0000016813AF58E66806C1E2D7F15155_1.csv",
      "size": 6705647390,
      "inventoriedRecord": 70585762,
    }
  ]
}
```

The name of a **manifest** file is as follows (for details about each field, see [Inventory File Name](#)):

```
destinationPrefix/sourceBucketName/inventoryId/yyyy-MM-dd'T'HH-mm'Z'/manifest.json
```

## How Is a User Notified When Inventory Files Are Generated?

You can enable SMN (the message notification service) for the destination bucket. By doing so, you can receive SMS messages or emails every time when inventory files and the **manifest** file are generated. For more information about SMN, see [Event Notifications](#).

The following is a simple example of SMN configuration. **destinationPrefix/sourceBucketName** indicates the prefix of the **manifest** file. **destinationPrefix** is the configured name prefix for inventory files, and **sourceBucketName** is the source bucket for which the inventory file is configured. The **manifest.json** is the suffix of the **manifest** file.

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Id>01</Id>
    <Filter>
      <Object>
        <FilterRule>
          <Name>prefix</Name>
          <Value>destination-prefix/source-bucket</Value>
        </FilterRule>
        <FilterRule>
          <Name>suffix</Name>
          <Value>manifest.json</Value>
        </FilterRule>
      </Object>
    </Filter>
  </TopicConfiguration>
</NotificationConfiguration>
```

```
<Event>ObjectCreated:Put</Event>
</TopicConfiguration>
</NotificationConfiguration>
```

## How to Use

You can configure bucket inventories using OBS Console.

Tool	Reference
OBS Console	<a href="#">Configuring a Bucket Inventory</a>

## 6.4 Event Notifications

You can configure to receive **Simple Message Notification (SMN)** notifications when certain events happen in your OBS bucket.

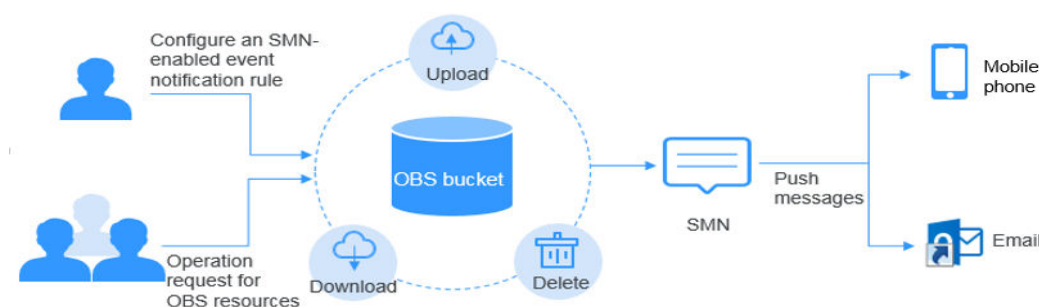
### SMN-Enabled Event Notification

Simple Message Notification (SMN) is a reliable and extensible message notification service that can handle a huge number of messages. SMN significantly simplifies system coupling. It can automatically push messages to subscribers through emails and text messages.

OBS uses SMN to provide the event notification function. You can use SMN to send event notifications to specified subscribers to inform them in real time of critical operations (such as upload and deletion) that occur on specified buckets. For example, you can configure an event notification rule to send messages through SMN to the specified email address whenever an upload operation occurs on the specified bucket.

You can configure an event notification rule to filter objects by the object name prefix or suffix. For example, you can add an event notification rule to send notifications whenever an object with the **.jpg** suffix is uploaded to the specified bucket. You can also add an event notification rule to send notifications whenever an object with the **images/** prefix is uploaded to the specified bucket.

**Figure 6-2** SMN-enabled event notification



## Event Types Supported by OBS

OBS can publish events of the following types. You need to specify these event types in the notification settings.

**Table 6-3** Event types supported by OBS

Event Type	Description
<b>ObjectCreated:*</b> (all upload operations) <b>ObjectCreated:Put</b> (object upload) <b>ObjectCreated:Post</b> (object upload through a web browser) <b>ObjectCreated:Copy</b> (copying objects) <b>ObjectCreated:CompleteMultipartUpload</b> (merging uploaded parts)	OBS can use APIs such as PUT, POST, and COPY to create objects and configure corresponding event types. You will receive a notification when an object is created using a specific API. In addition, you can use the <b>ObjectCreated:*</b> event type to request all object creation notifications.  <b>NOTE</b> You will not receive event notifications from failed operations.
<b>ObjectRemoved:*</b> (all deletion operations) <b>ObjectRemoved&gt;Delete</b> (deleting objects) <b>ObjectRemoved&gt;DeleteMarkerCreated</b> (A deletemarker object is created.)	By using the <b>ObjectRemoved</b> event types, you can enable notification when an object or a batch of objects are removed from a bucket. You can request notification when an object is deleted or a versioned object is permanently deleted by using the <b>ObjectRemoved&gt;Delete</b> event type. Alternatively, you can request notification when a delete marker is created for a versioned object by using <b>ObjectRemoved&gt;DeleteMarkerCreated</b> . You can also use <b>ObjectRemoved:*</b> to request notification each time an object is deleted.  <b>NOTE</b> You will not receive event notifications from automatic deletions of lifecycle policies or from failed operations.

## Constraints

- A configuration with no filtering attributes contained matches all prefixes and suffixes by default.
- When there are two suffixes configured, if a given string can end with both suffixes, this configuration is invalid due to suffix overlapping (for example, suffixes jpg and pg). This also works for prefixes.

## Destinations Supported by OBS

OBS can send event notification messages to **SMN** topics. You must grant OBS the permissions to send messages to these destinations. In addition, you need to specify the URN values of these destinations in the notification configuration.

## How to Enable Event Notifications

Enabling notifications is a bucket-level operation. OBS stores your event notification configuration to bucket sub-resources in the format of XML. By default, notification is not enabled for any event type. That is, the initial event notification configuration of each bucket is empty.

To enable notifications for events of specific types, you must add the corresponding XML configuration that identifies the event types you want OBS to publish and the destination where you want the notifications to be published.

- In this example, you need to publish event messages to an SMN topic. To set an SMN topic as the notification destination for specific event types, add the **TopicConfiguration**.

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Id>optional-id-string</Id>
    <Topic>topic-urn</Topic>
    <Event>event-type</Event>
    <Event>event-type</Event>
    ...
  </TopicConfiguration>
  ...
</NotificationConfiguration>
```

To remove all notification configurations from a bucket, set the **<NotificationConfiguration>** element to null.

## Using Object Key Names to Filter Event Notifications

You can configure event notifications based on the prefix and suffix of an object key name. For example, you can set up a configuration so that a notification is published only when objects with a **.jpg** extension are added to a bucket.

OBS stores the notification configuration as XML. You can use the **Filter** element in the XML structure to define the rules for notifications to be filtered by the prefix and/or suffix of an object key name. Notification configurations that use the filter cannot define filtering rules with overlapping prefixes, overlapping suffixes, or prefix and suffix overlapping. The following are examples of notification configurations with object key name filtering:

- **Example of valid notification configuration with object key name filtering**

The following information contains a configuration identifying an SMN topic to which OBS publishes events of the **ObjectCreated:Put** type. The events will be published each time an object that has an **image** prefix and a **jpg** suffix is PUT to a bucket.

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Id>01</Id>
    <Filter>
      <Object>
        <FilterRule>
```

```
<Name>prefix</Name>
<Value>image</Value>
</FilterRule>
<FilterRule>
  <Name>suffix</Name>
  <Value>jpg</Value>
</FilterRule>
</Object>
</Filter>
<Topic>urn:smn:southchina:11aa22bb:notification</Topic>
<Event>ObjectCreated:Put</Event>
</TopicConfiguration>
</NotificationConfiguration>
```

The following notification configuration has multiple non-overlapping prefixes. The configuration defines the following: When objects that have a prefix of **images** are uploaded to buckets, event notifications will be published to topic-A; when objects that have a prefix of **videos** are uploaded to buckets, event notifications will be published to topic-B.

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Id>01</Id>
    <Filter>
      <Object>
        <FilterRule>
          <Name>prefix</Name>
          <Value>images</Value>
        </FilterRule>
      </Object>
    </Filter>
    <Topic>urn:smn:southchina:11aa22bb:topic-A</Topic>
    <Event>ObjectCreated:Put</Event>
  </TopicConfiguration>
  <TopicConfiguration>
    <Id>02</Id>
    <Filter>
      <Object>
        <FilterRule>
          <Name>prefix</Name>
          <Value>videos</Value>
        </FilterRule>
      </Object>
    </Filter>
    <Topic>urn:smn:southchina:11aa22bb:topic-B</Topic>
    <Event>ObjectCreated:Put</Event>
  </TopicConfiguration>
</NotificationConfiguration>
```

The following notification configuration has multiple non-overlapping suffixes. The configuration defines the following: Notifications will be published to topic-A for all **.jpg** objects PUT to buckets, and notifications will be published to topic-B for all **.png** objects. The suffixes **.png** and **.jpg** are not overlapping even though they have the same last letter.

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Id>01</Id>
    <Filter>
      <Object>
        <FilterRule>
          <Name>suffix</Name>
          <Value>jpg</Value>
        </FilterRule>
      </Object>
    </Filter>
    <Topic>urn:smn:southchina:11aa22bb:topic-A</Topic>
    <Event>ObjectCreated:Put</Event>
  </TopicConfiguration>
```

```
<TopicConfiguration>
  <Id>02</Id>
  <Filter>
    <Object>
      <FilterRule>
        <Name>suffix</Name>
        <Value>.png</Value>
      </FilterRule>
    </Object>
  </Filter>
  <Topic>urn:smn:southchina:11aa22bb:topic-B</Topic>
  <Event>ObjectCreated:Put</Event>
</TopicConfiguration>
</NotificationConfiguration>
```

- **Examples of notification configuration with invalid prefix/suffix overlapping**

In most cases, notification configurations that use the filter cannot define filtering rules with overlapping prefixes, overlapping suffixes, or overlapping combinations of prefixes and suffixes for the same event types. (You can have overlapping prefixes if the suffixes do not overlap.) You can use overlapping object key name filtering rules with different event types. For example, you can create a notification configuration that uses the **image** prefix for the **ObjectCreated:Put** event type and the **image** prefix for the **ObjectDeleted:\*** event type.

A configuration with no filtering attributes contained matches all prefixes and suffixes by default. The following notification configuration is invalid because it contains overlapping prefixes. (The same thing would be true if **suffix** instead of **prefix** is used in this example.)

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Topic>urn:smn:southchina:11aa22bb:topic-A</Topic>
    <Event>ObjectCreated:*</Event>
  </TopicConfiguration>
  <TopicConfiguration>
    <Filter>
      <Object>
        <FilterRule>
          <Name>prefix</Name>
          <Value>abc</Value>
        </FilterRule>
      </Object>
    </Filter>
    <Topic>urn:smn:southchina:11aa22bb:topic-B</Topic>
    <Event>ObjectCreated:*</Event>
  </TopicConfiguration>
</NotificationConfiguration>
```

The following notification configuration is invalid because it contains overlapping suffixes. Two suffixes are considered overlapping if a given string can end with both suffixes. A string can end with **jpg** and **pg**. Therefore, the suffixes are overlapping. (The same is true for prefixes.)

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Filter>
      <Object>
        <FilterRule>
          <Name>suffix</Name>
          <Value>.jpg</Value>
        </FilterRule>
      </Object>
    </Filter>
    <Topic>urn:smn:southchina:11aa22bb:topic-A</Topic>
    <Event>ObjectCreated:*</Event>
```



```
</TopicConfiguration>
<TopicConfiguration>
  <Filter>
    <Object>
      <FilterRule>
        <Name>suffix</Name>
        <Value>pg</Value>
      </FilterRule>
    </Object>
  </Filter>
  <Topic>urn:smn:southchina:11aa22bb:topic-B</Topic>
  <Event>ObjectCreated:Put</Event>
</TopicConfiguration>
</NotificationConfiguration>
```

## Event Message Structure

A notification message sent by OBS to publish an event is a JSON message with the following structure.

```
{
  "Records": [
    {
      "eventVersion": "", //Version number. The current version is 3.0.
      "eventSource": "", //Message source. The value is fixed to OBS.
      "eventRegion": "", //Region where the event occurs
      "eventTime": "", //Time when an event occurs, in the ISO-8601 format, for example,
2020-07-10T09:24:11.418Z
      "eventName": "", //Name of the event that triggers the notification
      "userIdentity": {
        "ID": "" //Billing ID of the user who triggers the event
      },
      "requestParameters": {
        "sourceIPAddress": "" //Source IP address of the request
      },
      "responseElements": {
        "x-obs-request-id": "", //ID of the request
        "x-obs-id-2": "" //Special characters for locating problems
      },
      "obs": {
        "Version": "1.0",
        "configurationId": "", //Name of the event notification rule in OBS that matches the event
        "bucket": {
          "name": "examplebucket",
          "ownerIdentity": {
            "ID": "" //Account ID of the bucket owner
          },
          "bucket": "" //Bucket name
        },
        "object": {
          "key": "", //Object name
          "eTag": "", //ETag of the object
          "size": , //Object size
          "versionId": "null", //Version ID of the object
          "sequencer": "" //Identifier that defines the event sequence of a specific object
        }
      }
    }
  ]
}
```

Note the following:

- You can use the **responseElements** key-value to trace OBS requests. Both **x-obs-request-id** and **x-obs-id-2** can be used to trace a single request. Their value is the one returned by OBS in the response to the request.

- The **obs** key-value contains information about the bucket and object involved in the event. Note that the object key name is URL-encoded. For example, **TEST/Chinese.jpg** is changed to **TEST%2F%E4%B8%AD+%E6%96%87%2F.jpg**.

---

**NOTICE**

During secondary development, if you use the OBS SDK to download the object, you need to decode the URL-coded object name and then call the API for download. If a web browser is used to access the object, decoding is not required.

- The **sequencer** key-value determines the sequence of events. Generally, event notifications do not arrive in the order that the events occurred. However, notifications from events for creating objects (PUT) and deleting objects (DELETE) contain a **sequencer**, which can be used to determine the order of events for a given object key. If you compare the hexadecimal sequencer strings of two event notifications on the same object key, the event notification with the greater sequencer value is the event that occurred later.

---

**NOTICE**

1. The sequencer cannot be used to determine the sequence for events on different object keys.
2. The event sequence indicated by **sequencer** is for reference only, which does not work for highly reliable systems

---

Example message:

- Test message – When you configure an event notification on a bucket, OBS sends the following test message:

```
{
  "Service": "OBS",
  "Event": "TestEvent",
  "Time": "1970-01-01T00:00:00.000Z",
  "Bucket": "examplebucketname",
  "RequestId": "0002F4BCF6000001563B064B17B2094D",
  "HostId": "2Zf+b9AmbaBgNQ+YE8XU2j87DZaBNxu4TaMiOCTqpmkC2SA9ouf8TpB2SY5j3i4"
}
```

- Example message when an object is created using the PUT request – The following is an example of a message sent by OBS to publish an **ObjectCreated:Put** event:

```
{
  "Records": [
    {
      "eventVersion": "3.0",
      "eventSource": "OBS",
      "eventTime": "2018-06-26T14:37:05.468Z",
      "eventName": "ObjectCreated:Put",
      "userIdentity": {
        "ID": "71f3901173514e6988115ea2c26d1999"
      },
      "requestParameters": {
        "sourceIPAddress": "104.55.173.69"
      },
      "responseElements": {
```

```
"x-obs-request-id":"9006000001643C86D03C300BE8860FA7",
"x-obs-id-2":"2+/Ucr6uinCjAbUejWyQ+rhxkuf/K/9uoaXuewli/SE9j4tU5LwaXTTID1gvMv2o"
},
"obs":{
  "Version":"1.0",
  "configurationId":"ConfigurationId",
  "bucket":{
    "name":"examplebucket",
    "ownerIdentity":{
      "ID":"b4bf1b36d9ca43d984fbc9491b6fce9"
    },
    "bucket":"examplebucket"
  },
  "object":{
    "key":"object0001.txt",
    "eTag":"3b9680702b9a12733c5490d1b15c7607",
    "size":538,
    "versionId":"null",
    "sequencer":"000000001643C86D06576F5320000000"
  }
}
]
}
```

 **NOTE**

Due to Internet issues or constraints on e-mail sending, notifications may sometimes fail to reach HTTP or e-mail endpoints.

## How to Use

You can configure SMN notifications using OBS Console or APIs.

Tool	Reference
OBS Console	<a href="#">Configuring SMN Notifications</a>

# 7 Data Access

---

## 7.1 Static Website Hosting

### Precautions

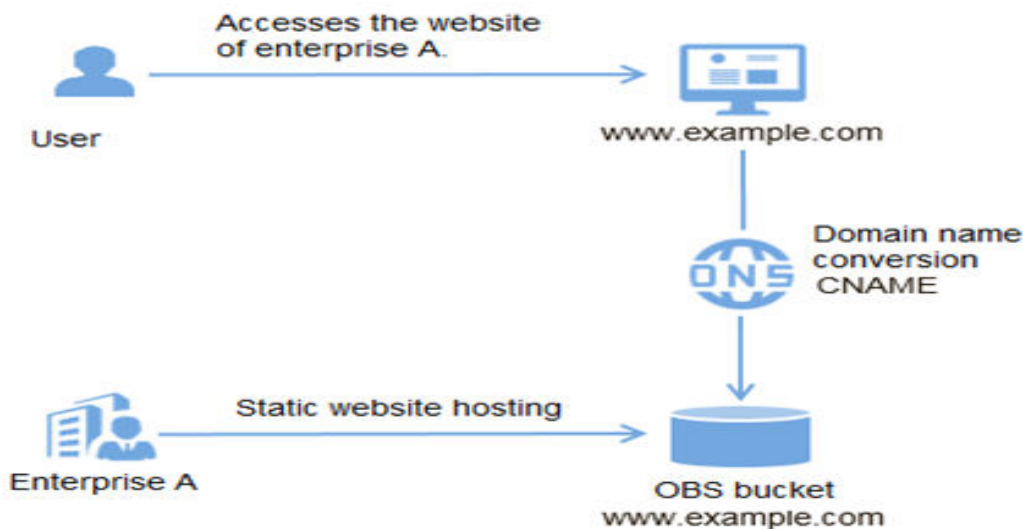
For security and compliance purposes, using static website hosting through the default OBS domain name (**bucket domain name or static website domain name**) will be prohibited by OBS. When you use such domain name to access web pages through a browser, no content will be displayed, instead, the content is downloaded as an attachment.

However, you can still use static website hosting through a user-defined domain name and can preview the web page content.

### Scenarios

You can use OBS to host your static website, for which you can also configure an index document, error document, and page direction. You can upload the content files of the static website to your bucket in OBS and configure a read permission to anonymous users for these files, and then configure the static website hosting mode for your bucket to host your static websites in OBS.

Static websites contain static web pages and some scripts that can run on clients, such as JavaScript and Flash. Different from static websites, dynamic websites rely on servers to process scripts, including PHP, JSP, and ASP.NET. OBS does not support scripts that run on servers.

**Figure 7-1** Static website hosting

After static website hosting is enabled, you can access objects in a bucket using either of the following methods:

- Resource management: You can access objects in a bucket through a common domain name. You can use APIs or SDKs and default OBS domain names with endpoints contained to perform common operations on buckets and objects, such as upload, download, deletion, and listing.
- Static website: You can access objects in a bucket by using a specific domain name. In this mode, display of the index page, error pages and requested pages is implemented according to the configured static domain name hosting rules.

The configuration of static website hosting takes effect within two minutes. After it takes effect, you can access static resources using the following URLs:

```
https://Static website domain name/Object name  
http://Static website domain name/Object name
```

A static website domain name is in the *BucketName.obs-website.Endpoint* format. *BucketName* indicates the name of a bucket, and *Endpoint* indicates the OBS domain name of the region where the bucket is located.

#### NOTE

Avoid using periods (.) in the destination bucket name. Otherwise, failures in client authentication certificate may occur when users use HTTPS for access.

To allow your clients to access the content on the website terminal node, you must make all of your content public and accessible. You can use bucket policies or ACLs on objects to grant permissions.

The following table lists the differences between the resource management mode and static website mode.

Major Difference	Resource Management	Static Website
Access control	Both public content and private content are supported.	Only public content is supported.
Error message processing	An error response in XML format is returned.	An HTML document is returned.
Redirection support	N/A	Both object-level and bucket-level redirection are supported.
Supported request	Operations on all buckets and objects are supported.	Only GET and HEAD requests on objects are supported.
Response to GET and HEAD requests of bucket root level	List of object keys in a bucket is returned.	Index file specified in the configuration is returned.

## Redirection Overview

When using static website hosting, you can configure redirection to redirect specific or all requests.

If the structure, address, or file name extension of a website is changed, users will fail to access the website using the old address (such as the address saved in the folder of favorites), and the 404 error message will be returned. In this case, you can configure redirection for the website to redirect user access requests to the specified page instead of returning the 404 error page.

Typical configurations include:

- Redirecting all requests to another website.
- Redirecting specific requests based on redirection rules.

## Configuring Static Website Hosting

### Overview

If you want to use a bucket to host static websites, add the website configuration to the bucket. The configuration includes the following information:

### Index document

When you enter a URL such as **http://example.com**, you are not requesting a specific page. In this case, the web server will provide a default page that contains the directory that stores the requested website content. This default page is called an index document, and in most cases it is named as **index.html**. When you configure a bucket for website hosting, you must specify an index document. When a request is sent to the root domain or any subfolder, OBS returns this index document.

### Error document

If an error occurs, OBS returns an HTML error document. For 4XX errors, you can provide your own customized error document, or provide other guides to your users in this document.

### Redirection of all requests

If the root domain is **example.com** and you need to respond to requests from **http://example.com** and **http://www.example.com**, you can create two buckets named **example.com** and **www.example.com**. Then you can retain the website content in only one bucket (such as **example.com**), and configure the other bucket to redirect all requests to the **example.com** bucket.

### Redirection based on advanced conditions

You can redirect requests based on the specific object name or prefix in the request, or based on the response code. For example, assume that you delete or rename an object in a bucket. You can add a routing rule that redirects requests to other objects.

The syntax format is as follows by setting redirection rules of specific requests:

```
<WebsiteConfiguration>
  <IndexDocument>
    <Suffix>index.html</Suffix></IndexDocument>
  <ErrorDocument>
    <Key>SomeErrorDocument.html</Key>
  </ErrorDocument>
  <RoutingRules>
    RoutingRules.....
  </RoutingRules>
</WebsiteConfiguration>
```

The syntax format of **RoutingRules** is as follows. The content in square brackets is optional.

```
<RoutingRules> =
  <RoutingRules>
    <RoutingRule>...</RoutingRule>
    [<RoutingRule>...</RoutingRule> ...]
  </RoutingRules>

<RoutingRule> =
  <RoutingRule>
    [ <Condition>...</Condition> ]
    <Redirect>...</Redirect>
  </RoutingRule>

<Condition> =
  <Condition>
    [ <KeyPrefixEquals>...</KeyPrefixEquals> ]
    [ <HttpErrorCodeReturnedEquals>...</HttpErrorCodeReturnedEquals> ]
  </Condition>

<Redirect> =
  <Redirect>
    [ <HostName>...</HostName> ]
    [ <Protocol>...</Protocol> ]
    [ <ReplaceKeyPrefixWith>...</ReplaceKeyPrefixWith> ]
    [ <ReplaceKeyWith>...</ReplaceKeyWith> ] [
    <HttpRedirectCode>...</HttpRedirectCode> ]
  </Redirect>
```

For request elements required for redirecting all requests sent to a specified website and for setting redirection rules, see [Configuring Static Website Hosting for a Bucket – Request Elements](#)

Examples:

### Example 1: Modifying the object name prefix for redirection

Assume that your bucket contains the following objects:

index.html

docs/article1.html

docs/article2.html

You decide to change the folder name from **docs/** to **documents/**. After the modification, the request for an object with prefix **/docs** needs to be redirected to another with **documents/**. For example, the request for **docs/article1.html** needs to be redirected to **documents/article1.html**.

In this case, you can add the following routing rules to the website configuration:

```
<RoutingRules>
  <RoutingRule>
    <Condition>
      <KeyPrefixEquals>docs/</KeyPrefixEquals>
    </Condition>
    <Redirect>
      <ReplaceKeyPrefixWith>documents/</ReplaceKeyPrefixWith>
    </Redirect>
  </RoutingRule>
</RoutingRules>
```

### Example 2: Redirecting requests sent to deleted folders to a specified page

Assume that you have deleted the **images/** folder, that is, you have deleted all objects whose object name prefix is **images/**. You can add a routing rule that redirects the requests of all objects whose prefix is **images/** to the page named **folderdeleted.html**.

```
<RoutingRules>
  <RoutingRule>
    <Condition>
      <KeyPrefixEquals>images/</KeyPrefixEquals>
    </Condition>
    <Redirect>
      <ReplaceKeyWith>folderdeleted.html</ReplaceKeyWith>
    </Redirect>
  </RoutingRule>
</RoutingRules>
```

### Example 3: Redirecting when HTTP errors occur

Assume that the requested object is not found, and the request needs to be redirected to **www.example.com**. You can add redirection rules so that site visitors can redirect to **www.example.com** when HTTP status code 404 (not found) is returned. The following example also inserts the object name prefix **report-404/** into redirection. For example, if you request the page **ExamplePage.html** and it causes an HTTP 404 error, the request will be redirected to the page **report-404/ExamplePage.html** on the **www.example.com**. If there is no routing rule and HTTP error 404 occurs, the error document specified in the configuration is returned.

```
<RoutingRules>
  <RoutingRule>
    <Condition>
      <HttpErrorCodeReturnedEquals>404</HttpErrorCodeReturnedEquals >
```



```
</Condition>
<Redirect>
  <HostName>www.example.com</HostName>
  <ReplaceKeyPrefixWith>report-404</ReplaceKeyPrefixWith>
</Redirect>
</RoutingRule>
</RoutingRules>
```

## Index Document Support

An index document is a web page returned when a request is sent to the root of a site or any subfolder. For example, if the user enters **http://www.example.com** in the browser, the user does not request any specific page. In this case, OBS provides index documents, which are sometimes referred to as default pages.

When configuring your bucket as a website, provide the name of the index document. You must upload the object with this name and configure it to be public and accessible.

Trailing slash in the root URL is optional. For example, if you configure an **index.html** website as an index document, **index.html** will be returned to any of the following URLs:

```
http://bucketname.obs-website.example.com/
http://bucketname.obs-website.example.com
```

In OBS, objects are horizontally stored in buckets. It does not provide any hierarchical organization as a file system on a computer. You can create a logical hierarchy by using the object name representing the folder structure. For example, consider a bucket with three objects and the following object names.

- sample1.jpg
- photos/2006/Jan/sample2.jpg
- photos/2006/Feb/sample3.jpg

Although they are not stored as any physical hierarchical organization, you can infer the following logical folder structure from the object name.

- The **sample1.jpg** object is located at the root level of the bucket.
- The **sample2.jpg** object is located in the **photos/2006/Jan** subfolder.
- The **sample3.jpg** object is located in the **photos/2006/Feb** subfolder.

## Customized Error Document Support

**Table 7-1** lists the subset of the HTTP response code returned by OBS when an error occurs.

**Table 7-1** List of OBS error codes

HTTP error code	Description
301 Moved Permanently	When a user sends a request to an endpoint of OBS, a <b>301 Moved Permanently</b> response is returned.

HTTP error code	Description
302 Found	When OBS receives a request for key <i>x</i> that does not end with a slash, it searches for the object whose key name is <i>x</i> . If the object is not found, OBS determines that the request is sent for subfolder <i>x</i> . Then OBS redirects the request by adding a slash at the end of the request, and returns <b>302 Found</b> .
304 Not Modified	OBS users request the <b>If-Modified-Since</b> and <b>If-Unmodified-Since</b> headers to determine whether the requested object is the same as the cached copy stored on the client. If the objects are the same, the website terminal node returns a <b>304 Not Modified</b> response.
403 Forbidden	When the user request is converted into an object that cannot be publicly read, the response of the website terminal node includes <b>403 Forbidden</b> . The object owner must use bucket policies or ACLs to make the object public and accessible.
404 Not Found	<p>The response of the website terminal node contains <b>404 Not Found</b> due to the following reasons:</p> <ul style="list-style-type: none"><li>• OBS determines that an invalid object key is referenced in the website URL.</li><li>• The OBS determines that the request is for an index document that does not exist.</li><li>• The bucket specified in the URL does not exist.</li><li>• The bucket specified in the URL exists but is not configured as a website.</li></ul> <p>You can create customized documents returned for <b>404 Not Found</b>. Ensure that the document has been uploaded to the bucket configured as a website and the website hosting configuration has been set to use the document.</p>

HTTP error code	Description
500 Service Error	When an internal server error occurs, the response of the website terminal node contains <b>500 Service Error</b> .
503 Service Unavailable	When OBS determines that you need to reduce the request frequency, the response of the website terminal node contains <b>503 Service Unavailable</b> .

## Permissions Required for Website Access

When you configure a bucket as a website, you must make the object to be provided public and accessible. To perform this operation, you need to write a bucket policy that grants GetObject permissions to everyone. If the object requested by the user does not exist on the terminal node of the website, OBS returns the HTTP response code **404 Not Found**. If the object exists but you do not grant read permission for the object, the website terminal node returns HTTP response code **403 Access Denied**. You can use this response code to infer whether a particular object exists. If you do not need this function, do not enable the bucket website support.

### NOTE

In the static website hosting scenario, anonymous users must be granted access to the hosted static website file. During their access to the hosted file, fees for outbound internet traffic and requests will be incurred.

The following example bucket policy grants each person the permission to access objects in a specified folder. For more information about bucket policies, see [Bucket Policies](#).

```
{
  "Statement": [{
    "Sid": "PublicReadGetObject",
    "Effect": "Allow",
    "Principal": {"ID": "*"},
    "Action": ["GetObject"],
    "Resource": ["example-bucket/*" ]
  }
  ]
}
```

### NOTICE

A bucket policy applies only to objects owned by a bucket owner. If the bucket contains objects that are not owned by the bucket owner, use object ACLs to grant the public read permission to the objects.

You can use bucket policies or object ACLs to grant public read permissions to your objects. To use ACLs to make objects public and accessible, you can grant the read permission to everyone, as shown in the following authorization elements.

You can add the authorization element to object ACLs. For details about ACL management, see [ACLs](#).

```
<Grant>
  <Grantee>
    <Canned>Everyone</Canned>
  </Grantee>
  <Permission>READ</Permission>
</Grant>
```

## Related Functions

Function	Relationship Description	Reference
CORS	By default, static websites hosted in an OBS bucket can only respond to access requests from websites in the same domain. You can configure cross-origin resource sharing (CORS) for the bucket so that the static websites can be accessed by cross-domain requests.	<a href="#">CORS</a>
URL Validation	OBS uses URL validation to prevent your websites hosted in an OBS bucket being stolen. OBS verifies URLs based on the referer field in the HTTP header.	<a href="#">URL Validation</a>

## How to Use

You can use OBS Console, APIs, or SDKs to configure static website hosting.

Tool	Reference
OBS Console	<a href="#">Configuring Static Website Hosting</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Configuring Static Website Hosting for a Bucket</a>

## 7.2 Domain Name Management

### 7.2.1 User-Defined Domain Name Configuration

#### Scenarios

If you want to migrate files from a website to OBS while keeping the website address unchanged, you can bind the website domain name to an OBS bucket, so that you can still use the website address to access the files stored in the bucket.

Assume the domain name of your website is **www.example.com** and the file you want to migrate to OBS is **abc.html**. After the file is migrated to OBS, you can use **http://www.example.com/abc.html** to access it. The steps below describe the configurations:

1. Create a bucket on OBS, and upload **abc.html** to the bucket.
2. On OBS Console, bind the domain name **www.example.com** to the created bucket.
3. Access the **abc.html** file. After the request for **http://www.example.com/abc.html** reaches OBS, OBS finds the mapping between the **www.example.com** and the bucket domain name, and redirects the request to the **abc.html** file stored in the bucket. The essence of this process is that OBS redirects the request to access **http://www.example.com/abc.html** to **http://bucket domain name/abc.html**.

## Constraints

### Bucket versions

- Only buckets with version 3.0 or later support user-defined domain name configuration. To check the bucket version, go to the **Overview** page of the bucket on OBS Console. Then you can view the bucket version in the **Basic Information** area.

### Number of domain names

- By default, a bucket can have up to 20 user-defined domain names bound.

### Functions

- User-defined domain names currently allow requests over only HTTP, but not HTTPS.

If you want to use a bound domain name to access OBS over HTTPS, you need to enable CDN to manage HTTPS certificates.

- A user-defined domain name can be bound to only one bucket.
- Chinese domain names are not supported.
- The suffix of a user-defined domain name can contain 2 to 6 uppercase or lowercase letters.

#### NOTE

If an acceleration domain name is also needed, to prevent objects in OBS buckets from being directly downloaded upon access, you need to perform other required operations after the custom domain name and the acceleration domain name have been configured. For details, see [With CDN Acceleration Enabled, Why Are the Objects in My OBS Bucket Directly Downloaded When I Access Them?](#)

## How to Use

You can configure user-defined domain name binding using OBS Console.

Tool	Reference
OBS Console	<a href="#">Configuring User-Defined Domain Names</a>

# 8 Data Security

---

## 8.1 Server-Side Encryption

### Scenarios

After server-side encryption is enabled, objects uploaded to OBS will be encrypted and then stored on the server. When objects are downloaded, they will be decrypted on the server first and then returned in plaintext to you.

OBS provides the following server-side encryption methods that adopt the 256-bit Advanced Encryption Standard (AES-256).

- Server-side encryption with keys hosted by KMS (SSE-KMS)

With this method, you need to create a key using Key Management Service (KMS) or use the default key provided by KMS. The KMS key is then used for server-side encryption when you upload objects to OBS.

You can enable SSE-KMS when creating a bucket. Then, all objects uploaded to the bucket can be encrypted. You can also enable SSE-KMS after a bucket is created. After SSE-KMS is enabled, the objects newly uploaded to the bucket will be encrypted.

OBS encrypts only the objects uploaded after the default encryption function is enabled. The encryption status of existing objects in the bucket remains unchanged. Disabling default encryption does not change the encryption status of existing objects in a bucket. After this function is disabled, you can still manually encrypt objects upon upload.

You can use OBS Console to configure SSE-KMS.

### Constraints

- Only one server-side encryption method can be used each time an object is uploaded.
- If SSE-KMS is enabled for a bucket or the objects in it, you must have the **kms:cmk:get**, **kms:cmk:list**, **kms:cmk:create**, **kms:dek:create**, and **kms:dek:crypto** permissions granted by using IAM, so that you can upload objects to or download objects from this the bucket.

## Background Information

In SSE-KMS mode, KMS uses a hardware security module (HSM) to protect key security, helping you easily create and control encryption keys. Keys are not displayed in plaintext outside HSMs, which prevents key disclosure. All operations performed on keys are controlled using access permissions and logged, meeting regulatory compliance requirements.

## Precautions

When server-side encryption is disabled for a bucket, the encrypted objects must be accessed over HTTPS.

## How to Use

You can use OBS Console to configure server-side encryption.

Tool	Reference
OBS Console	<a href="#">Uploading a File in Server-Side Encryption Mode</a> <a href="#">Configuring Bucket Default Encryption</a>

## 8.2 CORS

Cross-origin resource sharing (CORS) is a browser-standard mechanism provided by the World Wide Web Consortium (W3C). It defines the interaction methods between client-side web applications in one origin and resources in another. For general web page requests, website scripts and contents in one origin cannot interact with those in another because of Same Origin Policies (SOPs).

OBS supports CORS rules and allows resources in OBS to be accessed across origins. The configuration of CORS takes effect within two minutes.

OBS supports [Static Website Hosting](#). Static websites stored in OBS can respond to website requests from another origin only when CORS is configured for the bucket.

---

### NOTICE

By default, the OBS system is configured to support cross-domain access using the root domain name. This allows access from all domains, and clients are likely to be attacked.

To address this issue, you can create a **crossdomain.xml** file with specific rules in the bucket for each client, and add **Security.loadPolicyFile("https://bucket.obs.eu-west-101.myhuaweicloud.eu/crossdomain.xml")** in the flash code of the file to prevent attacks. **bucket.obs.eu-west-101.myhuaweicloud.eu** needs to be replaced with the actual access domain name of your bucket.

---

## Background Information

Cross-domain refers to access between different domains.

Restricting cross-domain access is a browser policy for security purposes, that is, the same-origin policy. Due to this JavaScript same-origin policy, JavaScript under domain A cannot operate objects under domain B or C.

The same protocol, domain name (or IP address), and port are considered as the same domain. If the protocols, domain names, and ports (if specified) of the two web pages are the same, the two web pages have the same origin. To better understand the same-origin policy, you can see the analysis on accessing the example address <https://support.huaweicloud.com/dir/test.html> in [Table 8-1](#).

**Table 8-1** Example analysis

URL	Access Result	Cause
<a href="https://support.huaweicloud.com/dir/other.html">https://support.huaweicloud.com/dir/other.html</a>	Successful	Same protocol, domain name, and port
<a href="https://support.huaweicloud.com/dir/inner/other.html">https://support.huaweicloud.com/dir/inner/other.html</a>	Successful	Same protocol, domain name, and port
<a href="http://support.huaweicloud.com/dir/test.html">http://support.huaweicloud.com/dir/test.html</a>	Failed	Same domain name and port, but different protocols
<a href="https://support.huaweicloud.com:81/dir/test.html">https://support.huaweicloud.com:81/dir/test.html</a>	Failed	Same protocol and domain name, but different ports
<a href="https://help.huaweicloud.com/dir/test.html">https://help.huaweicloud.com/dir/test.html</a>	Failed	Same protocol and port, but different domain names

## Scenarios

Typical application scenarios of CORS are as follows:

- Enables JavaScript and HTML5 to be used for establishing web applications that can directly access resources in OBS. No proxy servers are required for transfer.
- Enables the dragging function of HTML5 to be used to upload files to OBS (with the upload progress displayed) or update OBS contents using web applications.
- External web pages, style sheets, and HTML5 applications hosted in different origins can access web fonts or pictures stored in OBS, implementing resource sharing.

## How to Use

You can configure CORS using OBS Console, APIs, or SDKs.



Tool	Reference
OBS Console	<a href="#">Configuring CORS</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Configuring Bucket CORS</a>

## 8.3 Versioning

### Scenarios

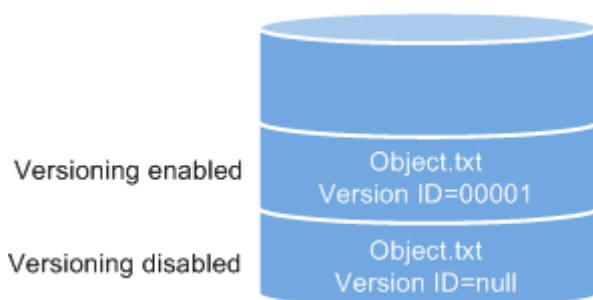
OBS can store multiple versions of an object. You can quickly search for and restore different versions or restore data in the event of accidental deletions or application faults.

By default, versioning is disabled for new OBS buckets. New objects will overwrite existing objects in case they have the same names.

### Enabling Versioning

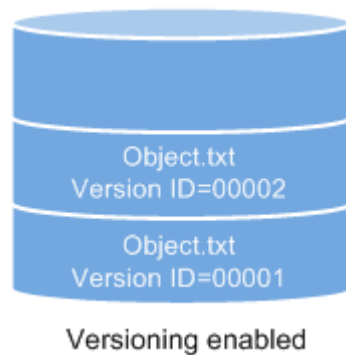
- Enabling versioning does not change the versions or contents of existing objects in the bucket. The version ID of an object is null before versioning is enabled. If a namesake object is uploaded after versioning is enabled, a version ID will be assigned to the object. For details, see [Figure 8-1](#).

**Figure 8-1** Versioning (enabled vs. disabled)



- With versioning enabled, OBS automatically allocates a unique version ID to a newly uploaded object. When an object with the same name as an existing object is uploaded again, both objects are stored in OBS with the same name but different version IDs. For details, see [Figure 8-2](#).

**Figure 8-2** Versioning (different version IDs for objects with the same name)



**Table 8-2** Version description

Version	Description
Latest version	After versioning is enabled, each operation on an object will result in saving of the object with a new version ID. The version ID generated upon the latest operation is called the latest version.
Historical Version	After versioning is enabled, each operation on an object will result in saving of the object with a new version ID. Version IDs generated upon operations other than the latest operation are called historical versions.

- The latest objects in a bucket are returned by default after a GET Object request.
- Objects can be downloaded by version IDs. By default, the latest object is downloaded if the version ID is not specified.
- You can select an object and click **Delete** on the right to delete the object. After the object is deleted, OBS generates a **Delete Marker** with a unique version ID for the deleted object, and the deleted object is displayed in the **Deleted Objects** list. If you try to access the deleted object, a 404 error will be returned.

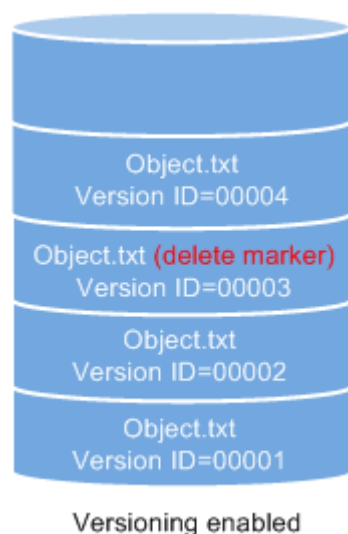
**Figure 8-3** Object with a delete marker



- You can recover a deleted object by deleting the version having a delete marker.
- After an object is deleted, you can specify the version number in **Deleted Objects** to permanently delete the object of the specified version.
- An object is displayed either in the object list or the list of deleted objects. It will never be displayed in both the lists at the same time.

For example, after object **A** is uploaded and deleted, it will be displayed in the **Deleted Objects** list. If you upload an object named **A** again, the object **A** will be displayed in the **Objects** list, and the previously deleted object **A** will no longer be displayed in the **Deleted Objects** list. For details, see [Figure 8-4](#).

**Figure 8-4** Uploading a namesake object after the original one is deleted

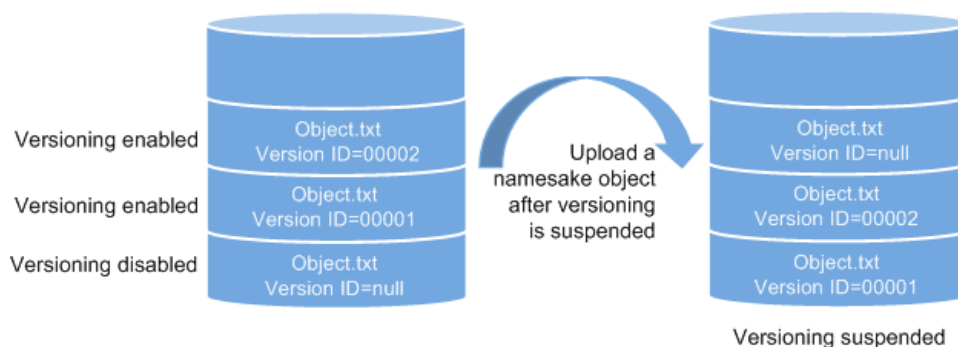


- All object versions except those with **Delete Marker** stored in OBS are charged.

## Suspending Versioning

Once the versioning function is enabled, it can be suspended but cannot be disabled. Once versioning is suspended, version IDs will no longer be allocated to newly uploaded objects. If an object with the same name already exists and does not have a version ID, the object will be overwritten.

**Figure 8-5** Object versions in the scenario when versioning is suspended



If versions of objects in a bucket do not need to be controlled, you can suspend the versioning function.

- Historical versions will be retained in OBS. If you do not need these historical versions, manually delete them.
- Objects can be downloaded by version IDs. By default, the latest object is downloaded if the version ID is not specified.
- All historical object versions except those with **Delete Marker** stored in OBS are charged.

## Differences Between Scenarios When Versioning Is Suspended and Disabled

If you delete an object when versioning is suspended, a null version with the **Delete Marker** is generated regardless of whether the object has historical versions. But, if versioning is disabled, the same operation will not generate a version with the **Delete Marker**.

### NOTE

After versioning is enabled, each historical version of an object is stored and occupies storage space. OBS charges storage fees for all versions. Exercise caution to avoid extra storage fees.

## How to Use

You can configure versioning through the OBS Console, APIs, and SDKs.

Tool	Reference
OBS Console	<a href="#">Configuring Versioning</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Configuring Versioning for a Bucket</a>

## 8.4 URL Validation

### Scenarios

Some rogue sites may steal links from other sites to enrich their content without any costs. Link stealing hurts the interests of the original websites and it is also a strain on their servers. OBS provides URL validation to solve this problem.

In HTTP, the **Referer** field allows websites and web servers to identify where people are visiting them from. URL validation of OBS utilizes this **Referer** field. The idea is that once you find that a request to your resource is not originated from an authorized source (for example, a URL), you can have the request blocked or redirected to a specific web page. This way, OBS prevents unauthorized access to data stored in buckets.

Such authorization is controlled using a whitelist and a blacklist.

**Referer setting rules:**

- The length of a whitelist or blacklist cannot exceed 1,024 characters.
- Referer format:
  - You can enter multiple referers, each in a line.
  - The referer parameter supports asterisks (\*) and question marks (?). An asterisk works as a wildcard that can replace zero or multiple characters, and a question mark (?) can replace a single character.
  - If the referer header field contains **http** or **https** during download, the referer must contain **http** or **https**.
- If there are websites configured in the blacklist referer, but no websites in the whitelist referer, all websites except those on the blacklist are allowed to access the target bucket.
- If there are websites configured in the whitelist referer, only the websites on the whitelist but not on the blacklist are allowed to access the target bucket, regardless of whether there are websites configured in the blacklist referer or not.

 **NOTE**

If a website is configured in both the whitelist and blacklist referers, access from this website will be forbidden. For example, if **https://www.example.com** is configured in both **Whitelisted Referers** and **Blacklisted Referers**, access requests from **https://www.example.com** will be blocked.

- If **Whitelisted Referers** and **Blacklisted Referers** are both left blank, all websites are allowed to access data in the target bucket by default.
- Before determining whether a user has the four types of permissions (**Read**, **Write**, **ACL View**, and **ACL Edit**) for a bucket or objects in the bucket, check whether this user complies with the URL validation principles of the **Referer** field.

**Whitelist and blacklist setting methods:**

- **Whitelist settings**

By setting a whitelist, you can allow requests from the websites in the whitelist, but deny those from the websites that are not in the whitelist.

For the requests that are initialized from a browser's address box, you can add the **\${null}** field to **Referer** of **Condition** to specify whether to allow the HTTP requests with a blank **referer**.

To configure a whitelist, refer to the following policy:

```
"Statement":[
  {"Sid": "1",
   "Effect": "Allow",
   "Principal": {"ID":["*"]},
   "Action": "*",
   "Resource":["bucket/*"],
  },
  {"Sid": "2",
   "Effect": "Deny",
   "Principal": {"ID":["*"]},
   "Action": ["*"],
   "Resource":["bucket/*"],
   "Condition":{
     "StringNotEquals":
     {"Referer":["http://www.example01.com","${null}"]}
   }
  }
]
```

```
}  
]
```

If you set a whitelist in this way, only the requests whose **referer** is set to **www.example01.com** or left blank can work on resources in buckets.

- Blacklist settings

To configure a blacklist, refer to the following policy:

```
"Statement":[  
  {"Sid":"1",  
   "Effect":"Deny",  
   "Principal":{"ID":["*"]},  
   "Action":["*"],  
   "Resource":["bucket/*"],  
   "Condition":{  
     "StringEquals":  
       {"Referer":["http://www.example01.com","http://www.example02.com"]}  
   }  
  }  
]
```

If you set a blacklist in this way, requests whose **referer** is set to **www.example01.com** or **www.example02.com** cannot work on resources in buckets.

## How to Use

You can use OBS Console or APIs to configure URL validation.

Tool	Reference
OBS Console	<a href="#">Configuring URL Validation</a>
API	<a href="#">Configuring a URL Validation Whitelist</a>

# 9 Monitoring and Logging

## 9.1 Monitoring

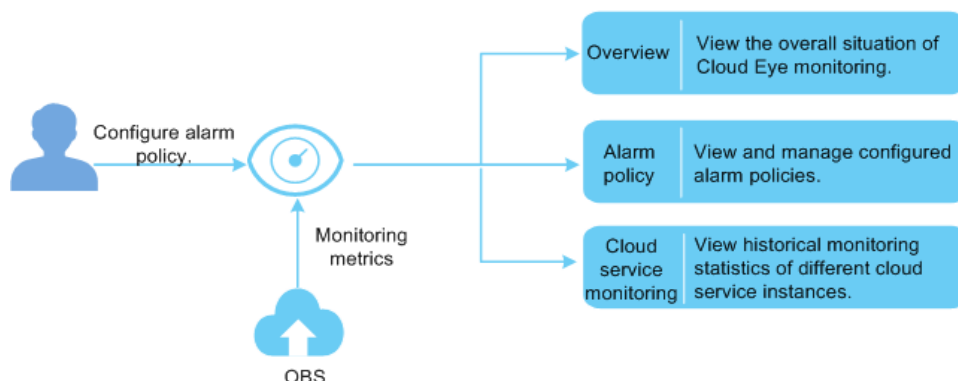
### Scenarios

You may send PUT and GET requests continuously when using OBS, which generates upload and download traffic. You may also receive error responses from the server. Huawei Cloud provides Cloud Eye to help you monitor OBS and better understand your bucket statuses. Cloud Eye can perform automatic and real-time monitoring over your buckets. It triggers alarms and notifications upon user operations based on predefined policies, allowing you to keep a close eye on your bucket access requests, traffic, and error responses.

You do not need to separately subscribe to Cloud Eye. It starts automatically once you create a resource (a bucket, for example) in OBS.

For more information about Cloud Eye, see [What Is Cloud Eye?](#)

**Figure 9-1** Cloud Eye monitoring



### Setting Alarm Rules

In addition to automatic and real-time monitoring, you can configure alarm rules in Cloud Eye to have alarm notifications sent when specific situations occur.

For details about how to configure alarm rules for Cloud Eye monitoring over OBS, see [Creating an Alarm Rule](#).

## Viewing OBS Monitoring Metrics

Cloud Eye monitors OBS usage in real time. You can view detailed statistics of each metric on the console of Cloud Eye.

For details about how to view OBS monitoring metrics, see [Querying Metrics of a Cloud Service](#).

## Monitoring Metrics

For details, see [OBS Monitoring Metrics](#).

# 9.2 Auditing

## Scenarios

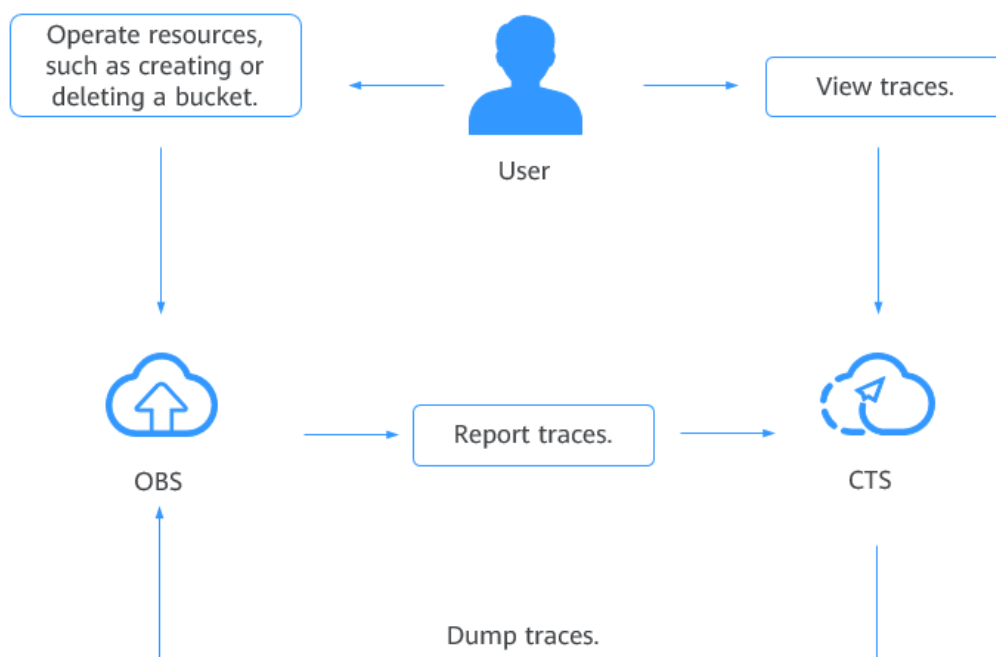
Cloud Trace Service (CTS) records operations on the cloud resources in your account. You can use the logs generated by CTS to perform security analysis, track resource changes, audit compliance, and locate faults.

After you enable CTS and configure a tracker, CTS can record management and data traces of OBS for auditing.

For details about how to enable and configure CTS, see [CTS](#).

For details about OBS management and data traces that can be traced by CTS, see [Cloud Trace Service](#).

**Figure 9-2** Cloud Trace Service





## 9.3 Logging

### Scenarios

You can enable logging to facilitate analysis or audit as required. Access logs enable a bucket owner to analyze the property, type, or trend of requests to the bucket in depth. With logging enabled, OBS automatically logs access requests for the bucket, and writes the generated log files to the specified bucket.

You need to specify a bucket for storing log files when enabling logging for a bucket. Log files can be stored in any bucket in the region where the logged bucket resides, including the logged bucket itself. To better manage logs, you are advised to store log files in a bucket other than the logged bucket. If log files are stored in the logged bucket, OBS creates additional logs for writing log files to the bucket, which takes up extra storage space that will increase your costs and makes it more difficult for you to locate required logs.

#### NOTICE

- Uploading bucket logs to the target bucket incurs billable PUT requests. For details about the pricing, see [Requests](#).
- After logging is configured for a bucket, you can view the bucket's operation logs in the log storage bucket in approximately 15 minutes.

OBS can record bucket access requests in logs for request analysis and log audit.

Logs occupy some OBS storage space rented by users, incurring extra fees. For this reason, the default policy is that OBS does not collect bucket access logs.

The log files are generated and uploaded by OBS to the bucket where the logs are stored. Therefore, OBS requires the authorization to upload the generated log files. Therefore, before configuring logging for a bucket, you need to create an IAM agency for OBS and add this IAM agency when configuring logging for the bucket. By default, when configuring permissions for an IAM agency, you only need to grant the IAM agency the permission to upload log files to the bucket where log files are stored. In the following example, **mybucketlogs** is the name of the bucket for storing log files. If the default encryption function is enabled for the log storing bucket, the IAM agency also requires the KMS Administrator permissions in the region where the log storing bucket resides.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "obs:object:PutObject"
      ],
      "Resource": [
        "OBS:*:*:object:mybucketlogs/*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

The following shows a sample access log record:

```
787f2f92b20943998a4fe2ab75eb09b8 bucket [13/Aug/2015:01:43:42 +0000] xx.xx.xx.xx
787f2f92b20943998a4fe2ab75eb09b8 281599BACAD9376ECE141B842B94535B
REST.GET.BUCKET.LOCATION
- "GET /bucket?location HTTP/1.1" 200 - 211 - 6 6 "-" "HttpClient" - -
```

The access log of each bucket contains the following information.

**Table 9-1** Format of bucket access logs

Name	Example	Description
BucketOwner	787f2f92b20943998a4fe2ab75eb09b8	Account ID of the bucket owner
Bucket	bucket	Name of the bucket
Time	[13/Aug/2015:01:43:42 +0000]	Timestamp of the request (UTC)
Remote IP	xx.xx.xx.xx	Request IP address
Requester	787f2f92b20943998a4fe2ab75eb09b8	Requester ID <ul style="list-style-type: none"> <li>When an account initiates a request, this parameter value is the account ID. When an IAM user initiates a request, this parameter value is the ID of the account where the IAM user belongs.</li> <li>When a request is initiated by an anonymous user, the value of this parameter is <b>Anonymous</b>.</li> </ul>
RequestID	281599BACAD9376ECE141B842B94535B	Request ID
Operation	REST.GET.BUCKET.LOCATION	Name of the operation
Key	-	Object name
Request-URI	GET /bucket?location HTTP/1.1	Request URI
HTTPStatus	200	Response code
ErrorCode	-	Error code
BytesSent	211	Size of the HTTP response, expressed in bytes

Name	Example	Description
ObjectSize	-	Object size (bytes) <b>NOTE</b> <ul style="list-style-type: none"> <li>When OBS deletes an object, it does not log the object's size. In the object deletion log, the value of <b>ObjectSize</b> is <b>0</b>.</li> <li>If error code <b>4XX</b> is returned, the value of <b>ObjectSize</b> is -, indicating that the specific object size is not displayed.</li> </ul>
TotalTime	6	Processing time on the server (ms)
Turn-AroundTime	6	Total time for processing the request (ms) <b>NOTE</b> This parameter can also be written as <b>TotalTime</b> .
Referer	-	Referer header of the request
User-Agent	HttpClient	User-Agent header of the request
VersionID	-	Version ID contained in a request
STSLogUrn	-	Federated authentication and agency information
StorageClass	STANDARD_IA	Current storage class of the object
TargetStorageClass	GLACIER	Storage class that the object will be transited to
DentryName	12456/file.txt	<ul style="list-style-type: none"> <li>For a parallel file system, this field indicates an internal identifier of a file or directory. Its value consists of a parent directory inode number and a file or directory name.</li> <li>For a bucket, the value of this field is -.</li> </ul>

## How to Use

You can configure logging on OBS Console, using APIs, or using SDKs.

Tool	Reference
OBS Console	<a href="#">Configuring Access Logging for a Bucket</a>
SDKs	OBS supports software development kits (SDKs) in multiple languages. For details, see the corresponding developer guide on the <a href="#">SDK Overview</a> page.
API	<a href="#">Configuring Logging for a Bucket</a>

# 10 Parallel File System

---

Parallel File System (PFS) is a high-performance file system provided by OBS, with access latency in milliseconds. PFS supports TB/s-level bandwidth and millions of IOPS, making it ideal for processing high-performance computing (HPC) workloads.

You can access data in a parallel file system using standard OBS APIs. It also supports data read and write through obsfs, a PFS client that supports POSIX. obsfs can be deployed on an ECS, and then you can use the obsfs to mount a parallel file system to the Linux OS running on the ECS. The mounted parallel file system functions as a local file system. You can manage the mounted parallel file system online, including creating, deleting, renaming files and folders, as well as modifying files.

For details, see [Parallel File System Feature Guide](#).

# 11 Change History

Released On	Description
2023-11-24	This is the eleventh official release. This issue incorporates the following change: <ul style="list-style-type: none"><li>Optimized the descriptions about percent-encoding for reserved characters in <a href="#">Object Overview</a>.</li></ul>
2023-11-03	This is the tenth official release. This issue incorporates the following change: <ul style="list-style-type: none"><li>Updated the description about <b>ObjectSize</b> in <a href="#">Table 9-1</a>.</li></ul>
2023-09-06	This is the ninth official release. <ul style="list-style-type: none"><li>Added the content about manually changing bucket and object storage classes in <a href="#">Storage Classes</a>.</li><li>Optimized the constraints.</li></ul>
2023-07-31	This is the eighth official release. <ul style="list-style-type: none"><li>Added <a href="#">Sharing Objects</a>.</li></ul>
2023-07-14	This is the seventh official release. <ul style="list-style-type: none"><li>Adjusted the document structure.</li></ul>
2023-06-16	This is the sixth official release. <ul style="list-style-type: none"><li>Added the rule naming precaution in <a href="#">Lifecycle Management</a>.</li></ul>
2023-06-07	This is the fifth official release. <ul style="list-style-type: none"><li>Added <a href="#">Direct Reading</a>.</li></ul>
2023-03-20	This is the fourth official release. <ul style="list-style-type: none"><li>Added <a href="#">Bucket Inventory</a>.</li></ul>
2023-03-15	This is the third official release. <ul style="list-style-type: none"><li>Added <a href="#">User-Defined Domain Name Configuration</a>.</li></ul>

Released On	Description
2023-03-14	This is the second official release. <ul style="list-style-type: none"><li>• Added <b>Server-Side Encryption</b>.</li><li>• Added <b>Tags</b>.</li></ul>
2022-09-30	This is the first official release.