ALIBABA CLOUD

Alibaba Cloud

Object Storage Service
Console User Guide

Document Version: 20220708

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
⚠ Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	⚠ Danger: Resetting will result in the loss of user configuration data.
🔔 Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	🔔 Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
🔈 Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	🔈 Notice: If the weight is set to 0, the server no longer receives new requests.
? Note	A note indicates supplemental instructions, best practices, tips, and other content.	? Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings> Network> Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.Overview 06

2.Manage buckets 07

2.1. Create buckets 07

2.2. Map custom domain names 13

2.3. Enable transfer acceleration 15

2.4. Access control 16

2.4.1. Configure hotlink protection for a bucket 16

2.4.2. Configure CORS 17

2.4.3. Modify the ACL of a bucket 19

2.5. Basic settings 21

2.5.1. Enable pay-by-requester 21

2.5.2. Configure bucket inventory 21

2.5.3. Configure server-side encryption 24

2.5.4. Configure bucket tagging 26

2.5.5. Configure static website hosting 26

2.5.6. Configure lifecycle rules 31

2.5.7. Configure retention policies 34

2.5.8. Back-to-origin rules 35

2.5.8.1. Overview 35

2.5.8.2. Basic configurations of mirroring-based back-to-orig... 36

2.5.8.3. Configure redirection-based back-to-origin 37

2.5.9. Configure event notification rules 39

2.5.10. Delete a bucket 41

2.6. Redundancy for fault tolerance 42

2.6.1. Configure CRR 42

2.6.2. Configure SRR 46

2.6.3. Configure versioning 48

2.7. Manage logs 50

2.7.1. Configure logging 51

2.7.2. Query real-time logs 51

2.8. Data Processing 54

2.8.1. Configure image styles 54

3.Upload, download, and manage objects 57

3.1. Overview 57

3.2. Upload objects 57

3.3. Create directories 60

3.4. Configure ACL for objects 61

3.5. Modify the storage class of objects 62

3.6. Configure bucket policies to authorize other users to acces.. 63

3.7. Download objects 70

3.8. Configure object metadata 71

3.9. Restore objects 73

3.10. Use OSS Select 74

3.11. Configure object tagging 75

3.12. Delete objects 76

3.13. Delete directories 77

4.Configure OSS DDoS protection 79

1.Overview

In the OSS console, you can perform basic and advanced operations on buckets, objects, and folders. The OSS console is a user-friendly and intuitive web application.

For more information about the OSS features and scenarios, see What is OSS?. For more information about the concepts, features, principles, and implementation methods, see OSS Developer Guide.

2.Manage buckets

2.1. Create buckets

A bucket is a container that is used to store objects in Object Storage Service (OSS). Before you upload an object to OSS, you must create a bucket.

Usage notes

- When you create a bucket, you are charged only for the storage of objects in the bucket and the traffic generated when the objects are accessed. For more information, see Overview.
- The capacity of the bucket is scalable. You do not need to purchase the capacity before you use the bucket.

Limits

- You can use an Alibaba Cloud account to create up to 100 buckets in the same region.
- A bucket name must be globally unique within OSS. For more information about the naming conventions of buckets, see Bucket naming conventions.
- After a bucket is created, its name, region, storage class, and redundancy type cannot be modified.
- OSS does not impose limits on the capacity of a bucket.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click **Create Bucket**.
3. In the **Create Bucket** panel, configure parameters described in the following table.

Parameter	Required	Description
Bucket Name	Yes	Specify the name of the bucket that you want to create. The name must meet the following requirements: ○ The bucket name must be globally unique in Alibaba Cloud OSS. ○ The name can contain only lowercase letters, digits, and hyphens (-). ○ The name must start and end with a lowercase letter or a digit. ○ The name must be 3 to 63 characters in length.

Parameter	Required	Description
Region	Yes	Select a region for the bucket. To access OSS from an Elastic Compute Service (ECS) instance over an internal network, select the region in which the ECS instance is located. For more information, see OSS domain names. ? **Note** You must complete real-name registration on the Real-name Registration page before you create a bucket in a region inside the Chinese mainland.

Parameter	Required	Description
Storage Class	Yes	Select a storage class for the bucket. ○ **Standard**: provides highly reliable, highly available, and high-performance object storage services that can handle frequent data access. Standard storage is ideal for storing images for social networking and sharing applications and storing data for audio and video applications, large websites, and big data analysis. ○ **IA**: provides high-durability storage services at a cost lower than Standard. Infrequent Access (IA) objects have a minimum storage period of 30 days and a minimum billable size of 64 KB. You can access IA objects in real time. However, you are charged data retrieval fees when you access IA objects. IA storage is suitable for data that is infrequently accessed, such as once or twice a month. ○ **Archive**: provides high-durability storage services at a cost lower than Standard and IA. Archive objects have a minimum storage period of 60 days and a minimum billable size of 64 KB. You must restore an Archive object before you can access it. The restoration takes approximately 1 minute. When you restore Archive objects, you are charged data retrieval fees. Archive storage is ideal for data that needs to be stored for a long period of time, such as archival data, medical images, scientific materials, and video footage. ○ **Cold Archive**: provides high-durability storage services at a cost lower than Standard, IA, and Archive. Cold Archive objects have a minimum storage period of 180 days and a minimum billable size of 64 KB. You must restore a Cold Archive object before you can access it. The amount of time required to restore a Cold Archive object depends on the object size and the restoration mode. When you restore Cold Archive objects, you are charged data retrieval fees. Cold Archive storage is ideal for storing cold data over an ultra-long period of time. Such data includes data that must be retained for an extended period of time due to compliance requirements, raw data that is accumulated over an extended period of time in the big data and AI fields, retained media resources in the film and television industries, and archived videos from the online education industry. **Note** Cold Archive storage is not supported only in the following regions: China (Nanjing - Local Region), South Korea (Seoul), and Thailand (Bangkok). For more information about storage classes, see Overview.

Parameter	Required	Description
OSS-HDFS	No	If you want to access OSS by using JindoSDK to build a data lake, enable the OSS-HDFS service. Before you enable OSS-HDFS, you must click **Authorize** and then follow the on-screen instructions in the panel to grant RAM users permissions to access OSS-HDFS. **Notice** ○ OSS-HDFS is supported only in the following regions: China (Hangzhou), China (Shanghai), China (Shenzhen), China (Beijing), China (Zhangjiakou), and Singapore (Singapore).To apply for a trial, contact technical support. OSS-HDFS cannot be disabled after it is enabled. Exercise caution when you enable OSS-HDFS. ○ OSS-HDFS cannot be enabled for Archive or Cold Archive buckets.
ZRS	No	Specify the redundancy type of the bucket. Valid values: ○ Activate: After this feature is enabled, OSS data is stored in zone-redundant storage (ZRS) mode. ZRS uses the multi-zone mechanism to distribute user data across three zones within the same region. Even if one zone becomes unavailable due to failures such as power outages and fires, the data is still accessible. **Notice** ZRS is supported only in the following regions: China (Shenzhen), China (Beijing), China (Hangzhou), China (Shanghai), China (Hong Kong), Singapore (Singapore), and Indonesia (Jakarta). You are charged extra fees for ZRS. This feature cannot be disabled after it is enabled. Exercise caution when you enable this feature. For more information about ZRS, see ZRS. ○ Not Activated: After ZRS is disabled, the redundancy type of the objects in the bucket is locally redundant storage (LRS). LRS stores the copies of each object across different devices within the same zone. This way, OSS ensures data reliability and availability even if two storage devices are damaged at the same time.

Parameter	Required	Description
Overview	No	Select whether to enable versioning. Valid values: ○ **Activate**: If you enable versioning for a bucket, objects that are overwritten or deleted in the bucket are stored as previous versions. Versioning allows you to recover objects in a bucket to a previous version, and protects your data from being accidentally overwritten or deleted. For more information, see Overview. ○ **Not Activated**: If you disable versioning for a bucket, objects that are overwritten or deleted in the bucket are not recovered.
ACL	Yes	Select the bucket ACL. Valid values: ○ **Private**: Only the bucket owner can perform read and write operations on objects in the bucket. Other users cannot access the objects in the bucket. ○ **Public Read**: Only the bucket owner can perform write operations on objects in the bucket. Other users, including anonymous users, can perform only read operations on the objects in the bucket. **Warning** All users on the Internet can access the objects in the bucket. This may result in unexpected access to the data in your bucket and unexpectedly high fees. Exercise caution when you set your bucket ACL to Public Read. ○ **Public Read/Write**: All users, including anonymous users, can perform read and write operations on the objects in the bucket. **Warning** All users on the Internet can access objects in the bucket and write data to the bucket. This may result in unexpected access to the data in your bucket and unexpectedly high fees. If a user uploads prohibited data or information, your legitimate interests and rights may be infringed. We recommend that you do not set your bucket ACL to Public Read/Write except in special cases.
Encryption Method	No	Select whether to enable server-side encryption for the bucket. **Note** Server-side encryption cannot be enabled only in the China (Nanjing - Local Region) region.

Parameter	Required	Description
Real-time Log Query	No	If you want to query OSS access logs of the last seven days free of charge, set Real-time Log Query to **Activate**. For more information about real-time log query, see Real-time log query. If you do not need to query real-time logs, keep the default setting, which is **Not Activated**.
Scheduled Backup	No	If you want to back up your OSS data on a regular basis, set Scheduled Backup to **Activate**. OSS automatically creates a backup plan to back up data once a day by using Hybrid Backup Recovery (HBR). The generated backup objects are stored for one week. **Notice** ◦ The scheduled backup feature is supported only in the following regions: China (Hangzhou), China (Shanghai), China (Shenzhen), China (Beijing), China (Zhangjiakou), China (Hong Kong), Singapore (Singapore), Australia (Sydney), Indonesia (Jakarta), and US (Silicon Valley). ◦ Scheduled backup cannot be configured for buckets whose storage classes are IA, Archive or Cold Archive. ◦ The backup and restoration of symbolic links, Archive and Cold Archive objects, and the access control lists (ACLs) of objects are not supported. ◦ If HBR is not activated or HBR is not authorized to access OSS, scheduled backup plans cannot be created. For more information, see Configure scheduled backup. If you do not need to back up your OSS data on a regular basis, keep the default setting, which is **Not Activated**.
Hierarchical Namespace	No	If you want to rename a directory or an object, enable the hierarchical namespace feature for the bucket in which the directory or object is stored. **Notice** You can enable the hierarchical namespace feature for a bucket only when you create the bucket. The hierarchical namespace feature cannot be disabled after it is enabled for a bucket. After you enable this feature for a bucket, some OSS features are no longer supported for the bucket. For more information about the features that are not supported for a bucket for which the hierarchical namespace feature is enabled, see Hierarchical namespace.

4. Click **OK**.

2.2. Map custom domain names

Prerequisites

- A custom domain name is registered. For more information about how to register a custom domain name, see Register a generic domain name.
- The Internet Content Provider (ICP) filing is complete for the domain name of the bucket that is located in the Chinese mainland and to which you want to map the custom domain name.

  For more information about how to apply for an ICP filing, see What is an ICP filing?.

Procedure

1. Map a custom domain name to a bucket.
   i. Log on to the OSS console.
   ii. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
   iii. In the left-side navigation pane, choose **Transmission > Domain Names**.
   iv. Click **Map Custom Domain Name**.
   v. In the **Map Custom Domain Name** panel, enter the domain name that you want to map in the **Custom Domain Name** field.

      Domain names that contain wildcard characters are not supported. Example: `*.example.com` .

      If a message appears, which indicates that a domain name conflict occurs, the domain name is mapped to another bucket. To resolve this issue, you can map another domain name to the bucket or verify the ownership of the domain name and forcibly map the domain name to the bucket. If you verify the ownership of the domain name and forcibly map the domain name to the bucket, the mapping between the domain name and another bucket is removed.
2. Add a CNAME record.
   - If the domain name is managed by your Alibaba Cloud account, perform the following steps to automatically add a CNAME record.
     a. In the **Map Custom Domain Name** panel, turn on **Add CNAME Record Automatically**.

        > **Notice** If the domain name has a CNAME record, the CNAME record is updated to the new CNAME record.

     b. Click **Submit**.
   - If the domain name is not managed by your Alibaba Cloud account, manually add a CNAME record.

     If the domain name is not hosted by Alibaba Cloud, you must add a CNAME record to the Domain Name System (DNS) of your DNS provider.

     The following example shows how to use Alibaba Cloud DNS to manually add a CNAME record for a domain name that is not hosted by Alibaba Cloud:

     a. Log on to the Alibaba Cloud DNS console.

b. On the Manage DNS page, click **Configure** in the Actions column of the domain name to which you want to add a CNAME record.

c. On the DNS Settings page, click **Add Record**. In the Add Record panel, configure the parameters. The following table describes the parameters.

Parameter	Description
Type	Select the type of the record that you want to add. In this example, **CNAME** is selected.
Host	Enter the host record based on the prefix of the domain name. ■ To add a top-level domain such as `aliyun.com`, enter **@**. ■ To add a second-level domain, enter the prefix of the second-level domain. If the domain name is `help.aliyun.com`, enter **help**. ■ To map all second-level domains to the domain name of the bucket, enter *****.
ISP Line	Select the ISP line that is used to resolve the domain name. To allow the system to select the optimal line, we recommend that you select **Default** for this parameter.
Value	Enter the domain name of the bucket. The domain name of a bucket is in the `BucketName.Endpoint` format. The public endpoint of the China (Hangzhou) region is `oss-cn-hangzhou.aliyuncs.com`. If you create a bucket named examplebucket in the China (Hangzhou) region, the domain name of the bucket is `examplebucket.oss-cn-hangzhou.aliyuncs.com`.
TTL	Select the interval at which the record is updated. In this example, the default value is used.

d. Click **Confirm**.

A new CNAME record immediately takes effect. The time required for a modified CNAME record to take effect is 72 hours.

Check whether the CNAME record takes effect

You can run the **ping** command or the **lookup** command to check whether a specified CNAME is in effect. If the request is redirected to `*.oss-cn-*.aliyuncs.com`, the CNAME is in effect.

Verify the ownership of a domain name

If a message appears, which indicates that a domain name conflict occurs, you can verify the ownership of the domain name and forcibly map the domain name to the bucket.

1. Click **Obtain TXT**.

   OSS randomly generates a token for the domain name, which includes the following fields: **Domain**, **Host**, and **Value**. Store the token information in a secure location.

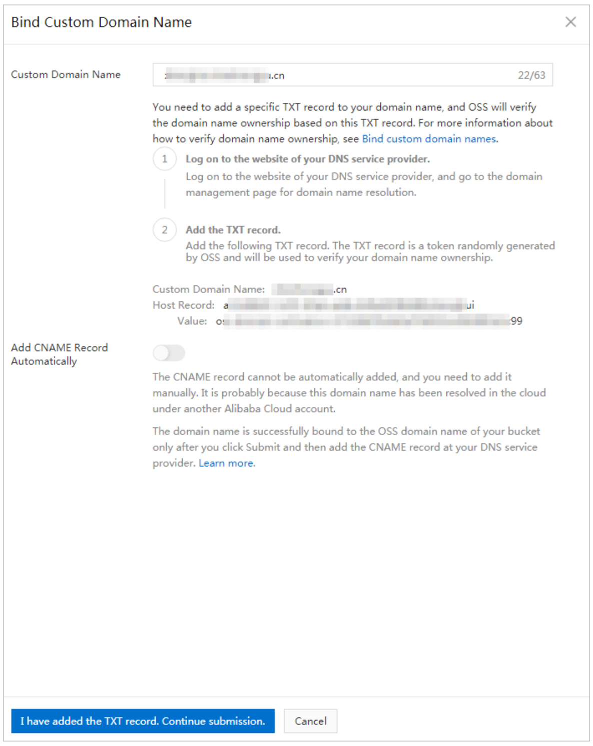

Bind Custom Domain Name

Custom Domain Name [redacted].cn 22/63

You need to add a specific TXT record to your domain name, and OSS will verify the domain name ownership based on this TXT record. For more information about how to verify domain name ownership, see Bind custom domain names.

1 Log on to the website of your DNS service provider.
Log on to the website of your DNS service provider, and go to the domain management page for domain name resolution.

2 Add the TXT record.
Add the following TXT record. The TXT record is a token randomly generated by OSS and will be used to verify your domain name ownership.

Custom Domain Name: [redacted].cn
Host Record: a[redacted]ui
Value: o[redacted]99

Add CNAME Record Automatically

The CNAME record cannot be automatically added, and you need to add it manually. It is probably because this domain name has been resolved in the cloud under another Alibaba Cloud account.

The domain name is successfully bound to the OSS domain name of your bucket only after you click Submit and then add the CNAME record at your DNS service provider. Learn more.

2. Add a TXT record to your DNS records. Enter the recorded values of the **Host** and **Value** fields and retain the default settings of other parameters.

   For more information about how to add a CNAME record, see Manually add a CNAME record.

3. In the **Map Custom Domain Name** panel, read and select **I have added the TXT record. Continue submission.**

   If your configurations are correct, OSS maps the custom domain name to the bucket.

Remove a domain name mapping

If you no longer need to use a custom domain name, you can remove the mapping of the custom domain name from a bucket.

1. On the management page of the bucket from which you want to remove the custom domain name mapping, choose **Transmission > Domain Names**.
2. On the Domain Names tab, click **Manage Mapping Configurations** in the Actions column of the domain name for which you want to remove the mapping.
3. In the **Manage Mapping Configurations** panel, click **Unbind**. Then, click **OK**.

2.3. Enable transfer acceleration

Object Storage Service (OSS) uses data centers distributed around the globe to implement transfer acceleration. When a request is sent to your bucket, it is parsed and routed to the data center where the bucket is located over the optimal network path and protocol. The transfer acceleration feature provides an optimized end-to-end acceleration solution to accessing OSS over the Internet.

Prerequisites

Real-name registration is complete.

You can complete real-name registration by submitting your information on the Real-name Registration page.

For more information, see Transfer acceleration.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Transmission > Transfer Acceleration**.
4. Click **Configure**, turn on Transfer Acceleration, and then click **Save**.

   Transfer acceleration takes effect within 30 minutes after it is enabled. After transfer acceleration is enabled for the bucket, the following two acceleration endpoints are added:

   - Global acceleration endpoint: *oss-accelerate.aliyuncs.com*. Transfer acceleration access points are distributed across the world. You can use this endpoint to accelerate data transfer for buckets in all regions.
   - Acceleration endpoint of regions outside the Chinese mainland: *oss-accelerate-overseas.aliyuncs.com*. Transfer acceleration access points are distributed across regions outside the Chinese mainland. You can use the acceleration endpoint to map a custom domain name without an ICP filing to a bucket in the China (Hong Kong) region or a region outside the Chinese mainland.

2.4. Access control

2.4.1. Configure hotlink protection for a bucket

You can enable the hotlink protection feature to configure a Referer whitelist for a bucket to prevent unauthorized access and associated unexpected fees.

Background information

Object Storage Service (OSS) provides the hotlink protection feature. This feature allows you to configure a Referer whitelist for a bucket. This way, only requests from domain names that are included in the Referer whitelist can access data in the bucket. You can configure Referer whitelists based on the Referer header field in HTTP and HTTPS requests.

After you configure a Referer whitelist for a bucket, OSS verifies requests to objects in the bucket only when the requests are initiated from signed URLs or anonymous users. Requests that contain the Authorization field in the header are not verified.

For more information about the API operation that you can call to configure a Referer whitelist for a bucket, see PutBucketReferer. For more information about hotlink protection, see Hotlink protection.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Access Control > Hotlink Protection**.
4. In the **hotlink protection** section, click **Configure**, and turn on **Enable Hotlink Protection**.

- Enter domain names or IP addresses in the **Referer Whitelist** field. Separate multiple Referers with new lines. You can use asterisks (*) and question marks (?) as wildcard characters. Examples:
  - If you add `www.aliyun.com` to the Referer whitelist, requests sent from URLs that start with *www.aliyun.com*, such as *www.aliyun.com/123* and *www.aliyun.com.cn* can access data in the bucket.
  - You can use an asterisk (*) as a wildcard character to specify zero or multiple characters. If you add `*www.aliyun.com/` to the Referer whitelist, requests sent from URLs such as *http://www.aliyun.com/* and *https://www.aliyun.com/* can access data in the bucket. If you add `*.aliyun.com` to the Referer whitelist, requests sent from URLs such as *help.aliyun.com* and *www.aliyun.com* can access data in the bucket.
  - You can use a question mark (?) as a wildcard character to specify a single character.
  - You can add domain names or IP addresses that include a port number, such as *www.example.com:8080* and *10.10.10.10:8080*, to the Referer whitelist.
- Select whether to turn on **Allow Empty Referer** to allow requests that contain an empty Referer.

  An HTTP or HTTPS request that contains an empty Referer indicates that the request does not contain the Referer field or that the Referer field is empty.

  If you turn off Allow Empty Referer, only HTTP or HTTPS requests that contain an allowed Referer field can access the objects in the bucket.

  > **Note** By default, if you preview an MP4 object by using a bucket domain name such as bucketname.oss-cn-zhangjiakou.aliyuncs.com, the browser sends two requests at the same time. One request contains the Referer field, and the other request contains an empty Referer. Therefore, you must add the bucket domain name to the Referer whitelist and turn on Allow Empty Referer. To preview a non-MP4 object by using the bucket domain name, you need to only allow requests that contain an empty Referer.

- Select whether to turn on **Truncate QueryString** to allow query strings to be truncated.

5. Click **Save**.

References

- If you want to allow users that meet specified conditions to access part of or all resources in your bucket or perform specific operations on the resources, we recommend that you configure bucket policies. For example, you can configure a bucket policy to allow only users from specified IP addresses or CIDR blocks to access a specified bucket. For more information about how to configure bucket policies, see Configure bucket policies to authorize other users to access OSS resources.
- For more information about how to troubleshoot hotlink protection errors, see Referer.

2.4.2. Configure CORS

Cross-origin resource sharing (CORS) is a standard cross-origin solution provided by HTML5 to allow web application servers to control cross-origin access. This way, the security of data transmission across origins is ensured.

Usage notes

- You can configure up to 10 CORS rules for a bucket.

- When Object Storage Service (OSS) receives a cross-origin request or an OPTIONS request that is destined for a bucket, OSS reads the CORS rules that are configured for the bucket and attempts to match the rules one after another. If OSS finds the first match, OSS returns corresponding headers. If the request fails to match the CORS rules, OSS does not include CORS headers in the response.
- To implement CORS after Alibaba Cloud CDN is activated, you must configure CORS rules in the CDN console. For more information, see Alibaba Cloud Content Delivery Network how to configure cross-origin resource sharing by using HTTP headers (CORS).

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Access Control > Cross-Origin Resource Sharing (CORS)**. In the **Cross-Origin Resource Sharing (CORS)** section, click **Configure**.
4. Click **Create Rule**. In the **Create Rule** panel, configure the parameters. The following table describes the parameters.

Parameter	Required	Description
Sources	Yes	The sources from which you want to allow cross-origin requests. When you configure the sources, take note of the following rules: ◦ You can configure multiple rules for sources. Separate multiple rules with line feeds. ◦ The domain names must include the protocol name, such as HTTP or HTTPS. ◦ You can use an asterisk (*) as the wildcard character. Each source can contain up to one asterisk (*). ◦ If a domain name does not use the default port, the domain name must contain the port number. Example: https://www.example.com:8080. The following examples show how to configure domain names: ◦ To match a specified domain name, enter the full domain name. Example: https://www.example.com. ◦ To match second-level domain names, use an asterisk (*) as the wildcard character in the domain name. Example: https://*.example.com. ◦ To match all domain names, enter only an asterisk (*) as the wildcard character.
Allowed Methods	Yes	The methods that cross-origin requests are allowed to use.

Parameter	Required	Description
Allowed Headers	No	The response headers for the allowed cross-origin requests. When you configure the headers, take note of the following rules: ○ This parameter is in the key:value format and not case-sensitive. Example: content-type:text/plain. ○ You can configure multiple response headers. Separate multiple response headers with line feeds. ○ Each rule can contain up to one asterisk (*) as the wildcard character. Set this parameter to an asterisk (*) if you do not have special requirements.
Exposed Headers	No	The response headers for allowed access requests from applications, such as an XMLHttpRequest object in JavaScript. Exposed headers cannot contain asterisks (*). We recommend that you set the following common exposed headers: ○ *x-oss-request-id* If you encounter an issue, contact technical support and provide the request ID to locate and resolve the issue. ○ *ETag* You can use the ETag value of an object to check whether the object content is modified.
Cache Timeout (Seconds)	No	The period of time in which the browser can cache the response to an OPTIONS preflight request for specific resources. Unit: seconds.
Vary: Origin	No	Specifies whether to return the Vary: Origin header. If both CORS and non-CORS requests are sent to OSS, or if the Origin header has multiple possible values, we recommend that you select the **Vary: Origin** header to avoid errors in the local cache. **Notice** If **Vary: Origin** is selected, visits through the browser or the CDN back-to-origin requests may increase.

For more information about the parameters, see PutBucketCors.

5. Click **OK**.

2.4.3. Modify the ACL of a bucket

Access control lists (ACLs) are used to control access to Object Storage Service (OSS) buckets and objects stored in OSS buckets. After a request is sent to access data stored in OSS, OSS checks the ACL of the data and verifies whether the requester has required permissions. You can configure the ACL of a bucket when you create the bucket. You can also modify the ACL of an existing bucket based on your requirements. Only the owner of a bucket can configure or modify the ACL of the bucket.

Usage notes

- If you do not specify the ACL of an object when you upload the object to a bucket, the ACL of the object inherits the ACL of the bucket.
- If you modify the ACL of a bucket, the ACLs of all objects that inherit the bucket ACL change accordingly.

ACL types

The following table describes bucket ACL types.

ACL	Operation
public-read-write	All users, including anonymous users, can perform read and write operations on the bucket. **Warning** When you set the bucket ACL to this value, all users can access the bucket and write data to the bucket over the Internet. This may result in unexpected access to the bucket and unexpectedly high fees. If a user uploads prohibited data or information, your legitimate interests and rights may be infringed. Therefore, we recommend that you do not set the bucket ACL to public-read-write except in special cases.
public-read	Only the owner of the bucket can write data to objects in the bucket. Other users, including anonymous users, can only read objects in the bucket. **Warning** When you set the bucket ACL to this value, all users can access objects in the bucket over the Internet. This may result in unexpected access to the bucket and unexpectedly high fees. Exercise caution when you set the bucket ACL to public-read.
private	Only the bucket owner can perform read and write operations on objects in the bucket. Other users cannot access the objects in the bucket. This is the default value.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Access Control > Access Control List (ACL)**.
4. In the **Access Control List (ACL)** section, click **Configure**. Then, modify the ACL of the bucket based on your requirements.
5. Click **Save**.

2.5. Basic settings

2.5.1. Enable pay-by-requester

You can enable the pay-by-requester mode for your bucket if you want requesters to pay the cost of requests and traffic generated when they access the data in your bucket. When pay-by-requester is enabled for a bucket, the requester instead of the bucket owner pays the cost of requests and traffic. The bucket owner pays only other fees including the storage fees.

Context

When the pay-by-requester mode is enabled, requesters pay for one or more of the following billable items based on their request content: number of API requests, outbound traffic over the Internet, back-to-origin traffic, Image Processing (IMG), video snapshots, and data retrieval of IA or Archive objects. The bucket owner pays other fees such as storage, object tagging, and transfer acceleration. For more information, see Enable pay-by-requester in OSS Developer Guide.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation bar, choose **Basic Settings > Pay by Requester**. In the **Pay by Requester** section, click **Configure** to enable or disable the pay-by-requester mode.
4. Click **Save**.

2.5.2. Configure bucket inventory

You can use the bucket inventory feature to export information about specified objects in a bucket, such as the number, sizes, storage classes, and encryption state of the objects. Compared with the GetBucket (ListObjects) operation, we recommend that you preferentially use the bucket inventory feature to list a large number of objects.

Usage notes

- You can configure up to 10 inventories for each bucket in the Object Storage Service (OSS) console and up to 1,000 inventories for each bucket by using OSS SDKs or ossutil.
- You are charged if you use the bucket inventory feature. However, only storage fees for inventory lists and API calling fees are charged during public preview.
- After an inventory is configured for a bucket, OSS generates inventory lists based on the inventory until the inventory is deleted. To save storage space, we recommend that you delete inventory lists that you no longer need in a timely manner.
- OSS keeps generating inventory lists at the interval specified by an inventory until the inventory is deleted. To prevent OSS from generating unnecessary inventory lists, you can delete inventories that you no longer need in a timely manner. You can also delete exported historical inventory lists that you no longer need.

For more information, see Bucket inventory.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Basic Settings > Bucket Inventory**. In the Bucket Inventory section, click **Configure**.
4. Click **Create Inventory**. In the **Create Inventory** panel, configure the parameters described in the following table.

Parameter	Description
Status	The status of the inventory. You can select **Enabled** or **Disabled**.
Rule Name	The name of the inventory. The name can contain only lowercase letters, digits, and hyphens (-) and cannot start or end with a hyphen (-).
Inventory Storage Bucket	The bucket in which generated inventory lists are stored. The source bucket for which an inventory is configured and the destination bucket in which the inventory lists are stored do not have to be the same bucket, but they must belong to the same account and reside within the same region.
Inventory List Path	The directory in which generated inventory lists are stored. ○ If you want to store the inventory lists in the root directory of the destination bucket, leave the parameter empty. ○ Otherwise, specify the parameter as a full path of a directory, excluding the destination bucket name. For example, if you want to store the inventory lists in the exampledir1 path of the destination bucket named examplebucket, set the parameter to *exampledir1*. If you want to store the inventory lists in the exampledir1/exampledir2 path of the destination bucket named examplebucket, set the parameter to *exampledir1/exampledir2*. **Note** If the path that you specify does not exist, OSS creates the path.
Frequency	The frequency at which inventory lists are generated. You can select **Weekly** or **Daily**. ○ If the number of objects in the bucket is up to 10 billion, set the parameter based on your business needs. ○ If the number of objects in the bucket is larger than 10 billion, we recommend that you set the parameter to Weekly.

Parameter	Description
Encryption Method	Specifies whether to encrypt inventory lists. ○ **None**: Inventory lists are not encrypted. ○ **AES-256**: Inventory lists are encrypted by using AES-256. ○ **KMS**: Inventory lists are encrypted by using a customer master key (CMK) managed by Key Management Service (KMS). To use a CMK to encrypt inventory lists, you must create a CMK in KMS in the same region as the destination bucket. For more information about how to configure CMKs, see Create a CMK. **Note** You are charged for calling API operations when you use CMKs to encrypt or decrypt data.
Object Versions	The object versions to which the inventory is applied. If versioning is enabled for the bucket, you can select **Current Version** or **All Versions**. For more information, see Overview. By default, inventory lists are generated for all objects in the bucket if versioning is not enabled for the bucket.
Object Prefix	The prefix based on which to scan objects. ○ If you want OSS to scan all objects in the bucket, leave the parameter empty. ○ To scan all objects in a path of the bucket, set this parameter to the full path that does not include the bucket name. For example, to scan all objects in the exampledir1 path of the eamplebucket bucket, set this parameter to *exampledir1/*. To scan all objects in the exampledir1/exampledir2 path of the examplebucket bucket, set this parameter to *exampledir1/exampledir2/*. **Note** If no objects in the bucket match the specified prefix, no inventory lists are generated.
Optional Fields	The object information that you want to include in inventory lists to be exported. You can select the following fields: **Object Size**, **Storage Class**, **Last Update Time**, **ETag**, **Multipart Upload**, and **Encryption Status**.

5. Read and select **I understand the terms and agree to authorize Alibaba Cloud OSS to access the resources in my buckets**. Then, click **OK**.
   It may take a large amount of time to generate inventory lists for a large number of objects. If you want to be notified when inventory lists are generated for the objects, we recommend that you configure an event notification for the destination bucket in which the inventory lists are stored and set the event to PutObject. When the inventory lists are generated, a notification is sent to you. For more information about how to configure event notifications, see Configure event notification rules.

2.5.3. Configure server-side encryption

Object Storage Service (OSS) allows you to configure server-side encryption. When you upload an object to a bucket for which server-side encryption is enabled, OSS encrypts the object before the object is stored. When you download the encrypted object from OSS, OSS automatically decrypts the object and returns the decrypted object. A header is added in the response to indicate that the object is encrypted on the OSS server.

Context

OSS provides the SSE-KMS and SSE-OSS methods for you to encrypt or decrypt data. The SSE-KMS method uses customer master keys (CMKs) stored in Key Management Service (KMS) to implement server-side encryption. SSE-OSS uses OSS-managed keys to implement server-side encryption. The following table describes the differences between the two methods and the scenarios of the two methods.

Encryption method	Description	Scenario	Billing
Server-side encryption by using SSE-KMS	You can use a default CMK or specify a CMK to encrypt or decrypt data. This method is cost-effective because you do not need to send user data to the KMS server over networks for encryption or decryption.	You can specify a customer-managed CMK to meet security and compliance requirements.	KMS charges you when you call API operations to encrypt or decrypt data by using CMKs stored in KMS.
Server-side encryption by using SSE-OSS	You can use SSE-OSS to encrypt each object. To improve security, OSS uses master keys that are rotated on a regular basis to encrypt data keys.	Only basic encryption capabilities are required. You do not need to manage keys on your own.	Free of charge

For more information about how the two encryption methods work and how you can implement the two encryption methods, see Server-side encryption.

You can enable server-side encryption in the OSS console by using one of the following methods:

- Method 1: Enable server-side encryption when you create a bucket
- Method 2: Enable server-side encryption for an existing bucket

Method 1: Enable server-side encryption when you create a bucket

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click **Create Bucket**.
3. In the **Create Bucket** panel, configure the parameters.

   Configure the following parameters in the **Server-side Encryption** section:

   - **Encryption Method**: Select an encryption method for the bucket.
     - **None**: Server-side encryption is disabled

    - **None**: Server-side encryption is disabled.
    - **OSS-Managed**: Keys managed by OSS are used to encrypt objects in the bucket. OSS uses data keys to encrypt objects. In addition, OSS uses regularly rotated master keys to encrypt data keys.
    - **KMS**: The default CMK stored in KMS or the specified CMK ID is used to encrypt and decrypt data.

      Before you use SSE-KMS, you must activate KMS. For more information, see activate KMS.
  - **Encryption Algorithm**:Only 256-bit Advanced Encryption Standard (AES-256) is supported.
  - **CMK**: You can set this parameter if you select **KMS** in the **Encryption Method** section. You can configure the following parameters for a CMK:
    - **alias/acs/oss**: The default CMK stored in KMS is used to encrypt different objects and decrypt the objects when they are downloaded.
    - CMK ID: The keys generated by a specified CMK are used to encrypt different objects, and the specified CMK ID is recorded in the metadata of the encrypted object. Objects are decrypted when they are downloaded by users who are granted decryption permissions. Before you specify a CMK ID, you must create a normal key or an external key in the same region as the bucket in the KMS console For more information, see Import key material.

   For other parameters, see Create buckets.

4. Click **OK**.

Method 2: Enable server-side encryption for an existing bucket

1. Log on to the OSS console.
2.
3. Choose **Basic Settings > Server-side Encryption**.
4. In the **Server-side Encryption** section, click **Configure**.

   You can configure the following parameters to enable server-side encryption:
   - **Encryption Method**: Select an encryption method for the bucket.
     - **None**: Server-side encryption is disabled.
     - **OSS-Managed**: Keys managed by OSS are used to encrypt objects in the bucket. OSS uses data keys to encrypt objects. In addition, OSS uses regularly rotated master keys to encrypt data keys.
     - **KMS**: The default CMK stored in KMS or the specified CMK ID is used to encrypt and decrypt data.

       Before you use SSE-KMS, you must activate KMS. For more information, see activate KMS.
   - **Encryption Algorithm**:Only 256-bit Advanced Encryption Standard (AES-256) is supported.
   - **CMK**: You can set this parameter if you select **KMS** in the **Encryption Method** section. You can configure the following parameters for a CMK:
     - **alias/acs/oss**: The default CMK stored in KMS is used to encrypt different objects and decrypt the objects when they are downloaded.

- CMK ID: The keys generated by a specified CMK are used to encrypt different objects, and the specified CMK ID is recorded in the metadata of the encrypted object. Objects are decrypted when they are downloaded by users who are granted decryption permissions. Before you specify a CMK ID, you must create a normal key or an external key in the same region as the bucket in the KMS console For more information, see Import key material.

5. Click **Save**.

> **Notice** The configurations of the default encryption method for a bucket do not affect the encryption configurations of existing objects within the bucket.

2.5.4. Configure bucket tagging

Object Storage Service (OSS) allows you to classify and manage buckets by using bucket tags. This topic describes how to configure bucket tagging by using the OSS console.

Context

For more information about bucket tagging, see Configure bucket tagging.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Basic Settings > Bucket Tagging** .
4. In the **Bucket Tagging** section, click **Configure**.
5. Enter the keys and values of tags that you want to configure for the bucket.

   To add multiple tags to the bucket, click the **+** icon.
6. Click **Save**.

2.5.5. Configure static website hosting

Static websites consist only of static content, including scripts such as JavaScript code that is run on the client. You can use the static website hosting feature to host your static website on an Object Storage Service (OSS) bucket and use the endpoint of the bucket to access the website.

Usage notes

For security reasons, starting from August 13, 2018 for regions inside mainland China, and September 25, 2019 for regions outside China, when you access web page objects whose MIME type is text/html and whose name extension is HTM, HTML, JSP, PLG, HTX, or STM by using browsers:

- If you use the default endpoint of the bucket to access the objects, the following header is automatically contained in the response: `Content-Disposition:'attachment=filename;'` . In this case, the web page objects are downloaded as attachments. The content of the object cannot be previewed.
- If you use a custom domain name mapped to the bucket to access the objects, the `Content-Disposition:'attachment=filename;'` header is not contained in the response. In this case, you can preview the object content if your browser supports preview of web page objects. For more

information about how to map a custom domain name to a bucket, see Map custom domain names.

For more information, see Overview.

Scenarios

In this topic, a bucket named examplebucket is used in the following example to show how to enable static website hosting for a bucket. The following structure shows the objects and directories in examplebucket:

```
examplebucket
  ├── index.html
  ├── error.html
  ├── example.txt
  └── subdir/
       └── index.html
```

The first example shows how to disable subdirectory homepage when you configure static website hosting for examplebucket. In this case, when you access the subdir/ directory of examplebucket, the default homepage object named index.html in the root directory of examplebucket is returned instead of the object named index.html in the subdir/ directory. In addition, if you access an object that does not exist in examplebucket, the specified default 404 page is returned. For more information, see Configure static website hosting and disable subdirectory homepage.

The second example shows how to enable subdirectory homepage and configure a subdirectory 404 rule when you configure static website hosting for examplebucket. In this case, when you access the subdir/ directory of examplebucket, the default homepage object named index.html in the /subdir directory of examplebucket is returned. In addition, if you access an object that does not exist in examplebucket, a result is returned based on the specified subdirectory 404 rule together with the specified default 404 page. For more information, see Configure static website hosting and enable subdirectory homepage.

Configure static website hosting and disable subdirectory homepage

1.
2.
3. In the left-side navigation pane, choose **Basic Settings > Static Pages**. Click **Configure** in the **Static Pages** section. Configure the parameters listed in the following table.

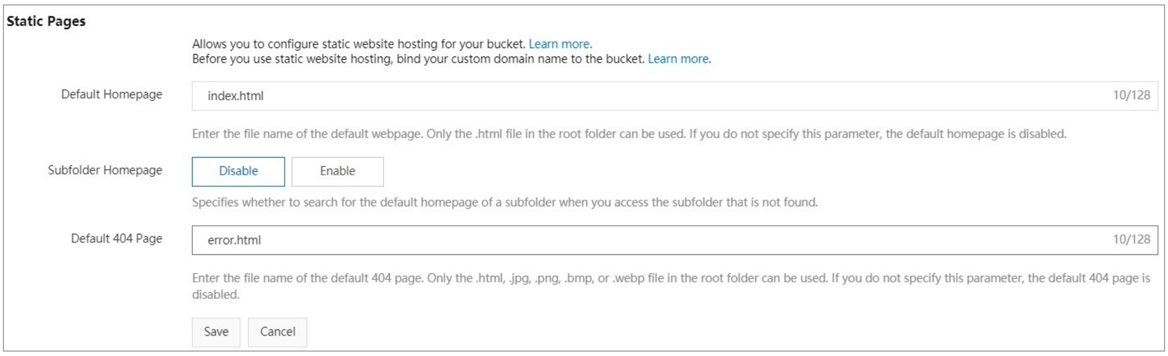

Static Pages

Allows you to configure static website hosting for your bucket. Learn more.
Before you use static website hosting, bind your custom domain name to the bucket. Learn more.

Default Homepage: index.html 10/128

Enter the file name of the default webpage. Only the .html file in the root folder can be used. If you do not specify this parameter, the default homepage is disabled.

Subfolder Homepage: [✓ Disable] Enable

Specifies whether to search for the default homepage of a subfolder when you access the subfolder that is not found.

Default 404 Page: error.html 10/128

Enter the file name of the default 404 page. Only the .html, .jpg, .png, .bmp, or .webp file in the root folder can be used. If you do not specify this parameter, the default 404 page is disabled.

Parameter	Description

Parameter	Description
Default Homepage	Configure the default homepage that appears when you use a browser to access the static website hosted on the OSS bucket. In this example, set this parameter to *index.html*.
Subfolder Homepage	Specify whether to enable subdirectory homepage for the bucket. In this example, select **Disable**. In this case, when you access the root directory of the bucket or a subdirectory whose URL ends with a forward slash (/), the default homepage object in the root directory of the bucket is returned.
Default 404 Page	Specify the error page that is returned when the object that you want to access does not exist in the bucket and a 404 HTTP status code is returned. Only an object in the root directory of the bucket can be specified as the default 404 page. In this example, set this parameter to *error.html*.
Error Page Status Code	Set the HTTP status code that is returned with the error page. Valid values: **404** and **200**.

4. Click **Save**.

Configure static website hosting and enable subdirectory homepage

1.
2.
3. In the left-side navigation pane, choose **Basic Settings > Static Pages**. In the **Static Pages** section, click **Configure**. Configure the parameters listed in the following table.

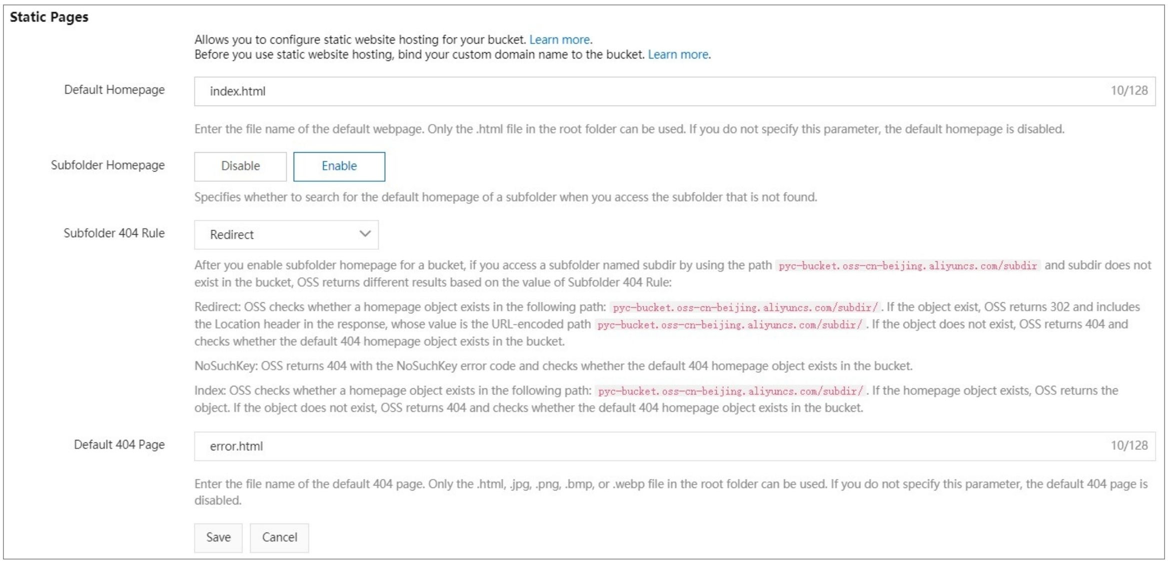

Static Pages

Allows you to configure static website hosting for your bucket. Learn more.
Before you use static website hosting, bind your custom domain name to the bucket. Learn more.

Default Homepage: index.html (10/128)

Enter the file name of the default webpage. Only the .html file in the root folder can be used. If you do not specify this parameter, the default homepage is disabled.

Subfolder Homepage: Disable [✓ Enable]

Specifies whether to search for the default homepage of a subfolder when you access the subfolder that is not found.

Subfolder 404 Rule: Redirect

After you enable subfolder homepage for a bucket, if you access a subfolder named subdir by using the path `pyc-bucket.oss-cn-beijing.aliyuncs.com/subdir` and subdir does not exist in the bucket, OSS returns different results based on the value of Subfolder 404 Rule:

Redirect: OSS checks whether a homepage object exists in the following path: `pyc-bucket.oss-cn-beijing.aliyuncs.com/subdir/`. If the object exist, OSS returns 302 and includes the Location header in the response, whose value is the URL-encoded path `pyc-bucket.oss-cn-beijing.aliyuncs.com/subdir/`. If the object does not exist, OSS returns 404 and checks whether the default 404 homepage object exists in the bucket.

NoSuchKey: OSS returns 404 with the NoSuchKey error code and checks whether the default 404 homepage object exists in the bucket.

Index: OSS checks whether a homepage object exists in the following path: `pyc-bucket.oss-cn-beijing.aliyuncs.com/subdir/`. If the homepage object exists, OSS returns the object. If the object does not exist, OSS returns 404 and checks whether the default 404 homepage object exists in the bucket.

Default 404 Page: error.html (10/128)

Enter the file name of the default 404 page. Only the .html, .jpg, .png, .bmp, or .webp file in the root folder can be used. If you do not specify this parameter, the default 404 page is disabled.

Parameter	Description
Default Homepage	Configure the default homepage that appears when you use a browser to access the static website hosted on the OSS bucket. In this example, set this parameter to *index.html*.

Parameter	Description
Subfolder Homepage	Select **Enable**. After you enable subdirectory homepage for a bucket, if you access the root directory of the bucket, the default homepage in the root directory is returned. If you access a subdirectory whose URL ends with a forward slash (/), the default homepage in the subdirectory is returned. For example, if you access `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/subdir/`, the default homepage object `index.html` in the *subdir/* directory is returned.
Subfolder 404 Rule	Configure the subdirectory 404 rule for the bucket. When you access an object that does not exist in the bucket, OSS returns different results based on the specified subdirectory 404 rule. For example, if you use the URL `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/exampledir` to access an object named *exampledir*, which does not exist in examplebucket, OSS returns different results based on the value that you set for this parameter. Default value: Redirect. ○ **Redirect**: OSS checks whether the *exampledir/index.html* object exists. ■ If this object exists, OSS returns 302 and redirects the request to `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/exampledir/index.html`. ■ If this object does not exist, OSS returns 404 and checks whether the `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/error.html` object exists. If this object also does not exist, OSS returns 404 status code. ○ **NoSuckKey**: OSS returns 404 and checks whether the `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/error.html` object exists. ○ **Index**: OSS checks whether the *exampledir/index.html* object exists. ■ If this object exists, OSS returns 200 and the content of this object. ■ If this object does not exist, OSS checks whether the `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/error.html` object exists.
Default 404 Page	Specify the error page that is returned when the object that you want to access does not exist in the bucket and a 404 HTTP status code is returned. Only an object in the root directory of the bucket can be specified as the default 404 page. In this example, set this parameter to *error.html*.
Error Page Status Code	Set the HTTP status code that is returned with the error page. Valid values: **404** and **200**.

4. Click **Save**.

Create and upload a default homepage object

If you set the default homepage to index.html when you configure static website hosting for examplebucket, you must upload an object named index.html to the root directory of examplebucket. If you enable subdirectory homepage for examplebucket, you must also upload index.html to the subdir/ directory of examplebucket.

1. Create a local file named index.html. The following example shows the content of index.html:

   ```
   <html>
   <head>
       <title>My Website Home Page</title>
       <meta charset="utf-8">
   </head>
   <body>
     <p>Now hosted on OSS.</p>
   </body>
   </html>
   ```

2. Save index.html to the local computer.
3. Upload index.html to the root directory and subdir/ directory of examplebucket. You must set the access control list (ACL) of index.html to public read.

   For more information about how to upload objects, see Upload objects.

Create and upload a default 404 page

If you set the default 404 page to error.html when you configure static website hosting for examplebucket, you must upload an object named error.html to the root directory of examplebucket.

1. Create a local file named error.html. The following example shows the content of error.html:

   ```
   <html>
   <head>
       <title>Hello OSS!</title>
       <meta charset="utf-8">
   </head>
   <body>
     <p>This is error 404 page.</p>
   </body>
   </html>
   ```

2. Save error.html to the local computer.
3. Upload error.html to the root directory of examplebucket. You must set the ACL of error.html to public read.

   For more information about how to upload objects, see Upload objects.

Disable static website hosting

If you no longer need to use the configurations of static website hosting, perform the following steps to disable the static website hosting feature:

1.
2.
3. In the left-side navigation pane, choose **Basic Settings > Static Pages**. In the **Static Pages**, click **Configure**.

4. Remove the configurations of the Default Homepage and Default 404 Page parameters and click **Save**.
   If a similar figure is returned, the static website hosting feature is disabled.

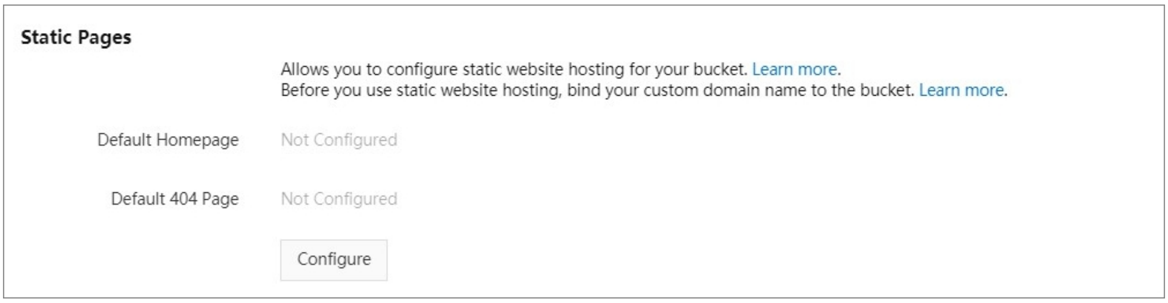

Static Pages

Allows you to configure static website hosting for your bucket. Learn more.
Before you use static website hosting, bind your custom domain name to the bucket. Learn more.

Default Homepage: Not Configured

Default 404 Page: Not Configured

2.5.6. Configure lifecycle rules

You can create lifecycle rules for a bucket based on the last modified time and last access time of objects in the bucket. This way, Object Storage Service (OSS) can regularly convert the storage class of the objects or delete expired objects and parts to reduce your storage costs.

Context

A lifecycle rule can contain policies based on both the last modified time and the last access time of objects. When you configure lifecycle rules, take note of the following limits:

- **Number of rules**

  You can configure up to 100 lifecycle rules in the OSS console. To configure more than 100 lifecycle rules, use OSS SDKs or the command-line tool ossutil.

- **Region**

  You can configure lifecycle rules based on the last access time only for buckets in the China (Qingdao), China (Hohhot), Germany (Frankfurt), and Australia (Sydney) regions.

- **Effective time**

  After a lifecycle rule is configured, the rule is loaded within 24 hours and takes effect within 24 hours after the rule is loaded. Check the configurations of a rule before you save the rule.

For more information about lifecycle rules that are configured based on the last modified time and the last access time of objects, see Overview.

Procedure

1.

2.

3. In the left-side navigation pane, choose **Basic Settings > Lifecycle**. In the **Lifecycle** section, click **Configure**.

4. Turn on **Enable access tracking** on the **Lifecycle** page if you want to create lifecycle rules based on the last access time of objects.

5. On the page that appears, click **Create Rule**. In the **Create Rule** panel, configure the options. The following table describes the options.

   ○ Options for unversioned buckets

Section	Parameter	Description
Basic Settings	**Status**	Specify the status of the lifecycle rule. Valid values: **Enabled** and **Disabled**.
	Applied To	Specify the objects to which the lifecycle rule applies. Valid values: **Files with Specified Prefix** and **Whole Bucket**.
	Allow Overlapped Prefixes	If you select this option, you can configure lifecycle rules with the same or overlapping prefixes without specifying tags. Notice - If you want OSS to automatically detect whether rules with the same or overlapping prefixes are configured, do not select this option. - If you want to use this option, contact technical support.
	Prefix	Specify the prefix in the names of objects to which the lifecycle rule applies. For example, if you want the rule to apply to objects whose names start with img, enter *img* in the field.
	Tagging	Specify tags. The rule applies only to objects that have the specified tags. For example, if you select **Files with Specified Prefix** and set Prefix to img, Key to a, and Value to 1, the rule applies to all objects that have the img prefix in their names and have the tag a=1. For more information about object tagging, see Object tagging.
	NOT	The NOT option is used to specify that the lifecycle rule does not apply to objects that have the specified prefix and tags. Notice - If you enable this option, each lifecycle rule must contain at least one of the prefix and tags of an object. - The key of the tag specified by the NOT syntax cannot be the same as the key specified by the **Tagging** option. - If you enable this option, lifecycle rules that apply to parts cannot be configured. - If you want to use this option, contact technical support.
	File Lifecycle	Configure rules for objects to specify when the objects expire. Valid values: **Validity Period (Days)**, **Expiration Date**, and **Disabled**. If you select **Disabled**, the configurations of File Lifecycle do not take effect.

Section	Parameter	Description
Policy for Objects		
	Lifecycle-based Rules	Configure the rule to convert the storage class of objects or delete expired objects. Example 1: If you select **Access Time**, set **Validity Period (Days)** to 30, and specify that the storage class of the objects is converted to **IA (Not Converted After Access)** after the validity period ends. In this case, the storage class of objects that were last accessed on September 1, 2021 is converted to Infrequent Access (IA) on October 1, 2021. **Note** If you configure a lifecycle rule based on the last access time, you can specify that the rule applies only to objects that are larger than 64 KB in size or to all objects in the bucket. Example 2: If you select **Modified Time**, set **Expiration Date** to September 24, 2021, and specify that objects that are last modified before this date are deleted. In this case, objects that are last modified before September 24, 2021 are automatically deleted. The deleted objects cannot be recovered.
Policy for Parts	**Part Lifecycle**	Specify the operations that you want to perform on expired parts. If you select **Tagging**, this option is unavailable. You can select **Validity Period (Days)**, **Expiration Date**, or **Disabled**. If you select **Disabled**, the configurations of Part Lifecycle do not take effect. **Notice** Each lifecycle rule must contain at least one of object expiration policies and part expiration policies.
	Rules for Parts	Specify when parts expire based on the value of Part Lifecycle. Expired parts are automatically deleted and cannot be recovered.

- Parameters for versioned buckets

  Configure the options in the **Basic Settings** and **Policy for Parts** sections in the same way as the options configured for unversioned buckets. The following table describes only the options that are different from those you configure for unversioned buckets.

Section	Parameter	Description
Policy for Current Versions	**Clean Up Delete Marker**	If you enable versioning for the bucket, you can configure the **Clean Up Delete Marker** option. Other options are the same as those you configure for unversioned buckets. If you select Clean Up Delete Marker, and an object has only one version which is a delete marker, OSS considers the delete marker expired and removes the delete marker. If an object has multiple versions and the current version of the object is a delete marker, OSS retains the delete marker. For more information about delete makers, see Delete marker.
Policy for Previous Versions	**File Lifecycle**	Specify when previous versions expire. Valid values: **Validity Period (Days)** and **Disabled**. If you select **Disabled**, the configurations of File Lifecycle do not take effect.
	Lifecycle-based Rules	Specify the number of days within which objects can be retained after they become previous versions. After they expire, the specified operations are performed on the previous versions the next day. For example, if you set Validity Period (Days) to 30, objects that become previous versions on September 1, 2021 are converted to the specified storage class or deleted on October 1, 2021. **Notice** You can determine when an object becomes a previous version based on the time when the next version of the object is last modified.

6. Click **OK**.

   After the lifecycle rule is saved, you can view the rule in the lifecycle rule list.

2.5.7. Configure retention policies

Object Storage Service (OSS) supports the Write Once Read Many (WORM) feature. The feature helps prevent objects from being deleted or overwritten within a specified period of time. Enterprises use this feature to comply with the regulations of the U.S. Securities and Exchange Commission (SEC) and Financial Industry Regulatory Authority (FINRA).

Usage notes

- You cannot configure retention policies in the China (Guangzhou), China (Nanjing - Local Region), and US (Virginia) regions.
- You can configure retention policies only for buckets in OSS.
- A bucket cannot have versioning and retention policies at the same time. If versioning is enabled for a bucket, you cannot configure retention policies for the bucket. For more information about versioning, see Overview.
-
- During the retention period, you can configure lifecycle rules to convert the storage classes of the objects in the bucket. This way, you can reduce costs and ensure compliance. For more information

about lifecycle rules, see Overview.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which you want to create a directory.
3. In the left-side navigation pane, choose **Basic Settings > Retention Policy**. In the **Retention Policy** section, click **Configure**.
4. Click **Create Policy**.
5. In the **Create Policy** dialog box, set **Retention Period**.

   The retention period ranges from one day to 70 years.
6. Click **OK**.

   After you create the policy, the policy is in the InProgress state. You can click **Lock** or **Delete** to lock or delete a policy in the InProgress state.
7. Click **Lock**.
8. In the message that appears, click **OK**.

   > Notice
   >
   > - The policy enters the Locked state. You cannot delete the policy or shorten the retention period. However, you can click **Edit** to extend the retention period.
   > - During the retention period, data in the bucket is protected. If you attempt to delete or modify the data, the following error message is displayed: `The file is locked and cannot be operated`.

Calculate the expiration time of an object

To calculate the time when an object expires, add the retention period and the time when the object was last updated. For example, the retention policy for Bucket A specifies the retention period as 10 days. An object in the bucket was last updated at 12:00 on February 15, 2022. The object expired at 12:01 on February 25, 2022.

2.5.8. Back-to-origin rules

2.5.8.1. Overview

You can configure back-to-origin rules for a bucket. If a request is sent to access an object but the object does not exist in a bucket for which back-to-origin rules are configured, Object Storage Service (OSS) obtains the requested object from the origin specified by the back-to-origin rules. You can configure mirroring-based or redirection-based back-to-origin rules for hot migration and specific request redirection.

For more information, see Manage back-to-origin configurations.

Mirroring-based back-to-origin

After you configure mirroring-based back-to-origin rules for a bucket, if a requester accesses an object but the object does not exist in a bucket, OSS obtains the object from the origin specified by the back-to-origin rules. OSS returns the object obtained from the origin to the requester and stores the object in the bucket. For more information about how to configure mirroring-based back-to-origin rules, see Basic configurations of mirroring-based back-to-origin.

Mirroring-based back-to-origin is used to migrate data to OSS. This feature allows you to migrate a service that already runs on an origin that you create or on another cloud service to OSS without service interruptions. You can use mirroring-based back-to-origin rules during migration to obtain the data that is not migrated to OSS. This ensures service continuity. For detailed examples, see Seamlessly migrate data of a web-based service provider to OSS.

Redirection-based back-to-origin

After you configure redirection-based back-to-origin rules for a bucket, if an error occurs when a requester accesses the bucket, OSS redirects the request to the origin specified by the redirection-based back-to-origin rules. You can use this feature to redirect requests for objects and develop various services based on redirection. For more information about how to configure redirection-based back-to-origin rules, see Configure redirection-based back-to-origin.

Rules

You can configure up to 20 back-to-origin rules for a bucket in the OSS console. By default, the rules are used to match a request in the sequence that they are configured. You can click **Move Up** or **Move Down** on the right side of the rules to change the priority of the rules.

> **Notice** If a request matches a rule, subsequent rules are not used to match the request. OSS determines whether a request matches a back-to-origin rule based on whether the request meets the conditions specified in the rule. OSS does not check whether a request can obtain the requested object from the origin.

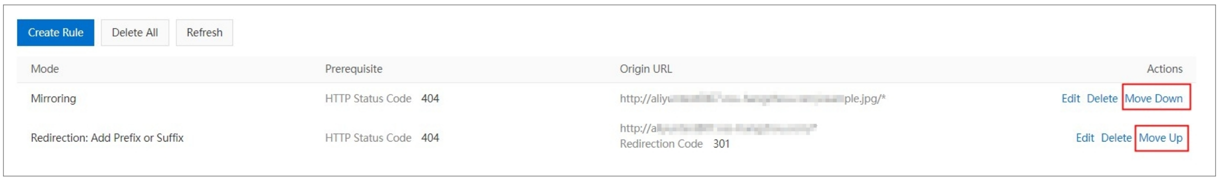

Mode	Prerequisite	Origin URL	Actions
Mirroring	HTTP Status Code 404	http://aliy[redacted]ple.jpg/*	Edit Delete Move Down
Redirection: Add Prefix or Suffix	HTTP Status Code 404	http://a[redacted] Redirection Code 301	Edit Delete Move Up

2.5.8.2. Basic configurations of mirroring-based back-to-origin

Mirroring-based back-to-origin is used to seamlessly migrate data to OSS. This feature allows you to migrate a service that already runs on an origin that you create or on another cloud service to OSS without service interruptions. You can use mirroring-based back-to-origin rules during migration to obtain the data that is not migrated to OSS. This ensures service continuity.

Procedure

When a requester accesses an object in the specified bucket but the object does not exist, you can specify the URL of the object in the origin and back-to-origin conditions to obtain the object. For example, you have a bucket named example that is located in the China (Hangzhou) region. You want your requester to obtain the required object in the *examplefolder* folder from `https://www.example.com/` when the requester accesses the object that does not exist in the *examplefolder* folder of the bucket root folder. To configure a back-to-origin rule, perform the following steps:

1.
2.
3. In the left-side navigation pane, choose **Basic Settings > Back-to-Origin**.
4. Click **Configure**. Click **Create Rule**.
5. In the **Create Rule** panel, set **Mode** to **Mirroring**.
6. Configure **Prerequisites** and **Origin URL**.

Parameter	Description
Prerequisite	Select **File Name Prefix**, and then set File Name Prefix to **examplefolder/**. **Note** File Name Prefix and File Name Suffix are optional. When you configure multiple back-to-origin rules, you must set different file name prefixes or suffixes to differentiate back-to-origin rules.
Origin URL	Set the first column to **https**, the second column to **www.example.com**, and the third column to **examplefolder**.

7. Click **OK**.

   Access process after the back-to-origin rule is configured:

   i. A requester accesses `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/examplefolder/example.txt` for the first time.

   ii. If the *examplefolder/example.txt* object does not exist in examplebucket, OSS sends a request to `https://www.example.com/examplefolder/example.txt` .

   iii. If the required object is obtained, OSS writes the *example.txt* object to the *examplefolder* folder in examplebucket and then returns the object to the requester. If the required object is not obtained, HTTP status code 404 is returned to the requester.

References

For more information about mirroring-based back-to-origin and some special scenarios, see Special configurations of mirroring-based back-to-origin.

2.5.8.3. Configure redirection-based back-to-origin

After you configure redirection-based back-to-origin rules for a bucket, if an error occurs when a requester accesses the bucket, OSS redirects the request to the origin specified by the redirection-based back-to-origin rules. You can use this feature to redirect requests for objects and develop various services based on redirection.

Procedure

If an error occurs when your visitor accesses a bucket, you can specify the URL of the object in the origin and back-to-origin conditions to redirect the access to the origin. For example, you have a bucket named example that is located in the China (Hangzhou) region. You want a requester to be redirected to obtain the required object in the *examplefolder* folder from `https://www.example.com/` when the requester accesses the object in the *examplefolder* folder of the bucket root folder but the object does not exist.

1.
2.
3. In the left-side navigation pane, choose **Basic Settings > Back-to-Origin**.
4. Click **Configure**. Click **Create Rule**.
5. In the **Create Rule** panel, set **Mode** to **Redirection**.
6. Configure **Prerequisite** and **Origin URL**.

Parameter	Description
Prerequisite	○ Select **HTTP Status Code**, and then set HTTP Status Code to **404**. Valid values of HTTP status codes: 400 to 599. For more information about the error information of various HTTP status codes, see Common errors and troubleshooting. ○ Select **File Name Prefix**, and then set File Name Prefix to **examplefolder/**. **Note** File Name Prefix and File Name Suffix are optional. When you configure multiple back-to-origin rules, you must set different file name prefixes or suffixes to differentiate back-to-origin rules.
Origin URL	Select **Add Prefix or Suffix**. Set the first column to **https** and the second column to **www.example.com**. Leave another column empty.

7. Click **OK**.

   Access process after the back-to-origin rule is configured:

   i. A requester accesses `https://examplebucket.oss-cn-hangzhou.aliyuncs.com/examplefolder/example.txt` for the first time.

   ii. If the *examplefolder/example.txt* object does not exist in examplebucket , OSS returns HTTP status code 301 to the requester and provides `https://www.example.com/examplefolder/example.txt` for redirection.

iii. The requester accesses `https://www.example.com/examplefolder/example.txt`.

The following table describes the parameters you can configure for different scenarios.

Scenario	Parameter
An object name prefix is different from the prefix of the name of the object in the origin.	Select **Replace or Delete File Prefix**, and set the third column of **Origin URL**. OSS replaces the **File Name Prefix** content with the content in the third column of **Origin URL**. You can configure this item after you configure **File Name Prefix**.
Transfer the query string included in a request from OSS to the origin.	Select **Transfer queryString**.
An HTTP redirect code must be replaced.	The default HTTP redirect code of a redirection rule is 301. You can select **302** or **307** from the **Redirection Code** drop-down list.
A redirect request is from Alibaba Cloud CDN.	Specify whether to select **Source from Alibaba Cloud CDN**. When the redirect request is from Alibaba Cloud CDN: If you select **Source from Alibaba Cloud CDN**, CDN automatically follows redirection rules and then obtains content. If you do not select **Source from Alibaba Cloud CDN**, CDN automatically returns the URL for redirection to the client.

2.5.9. Configure event notification rules

You can configure event notification rules for objects that you want to monitor in the Object Storage Service (OSS) console. If the events that you specify in the rules occur on these objects, you can receive notifications from the specified HTTP servers or Message Service (MNS) queues in real time.

Prerequisites

MNS is activated. You can go to the MNS product page to activate MNS.

Usage notes

- You are charged by MNS when you use the event notification feature. For more information about the pricing, see Pricing.
- The event notification feature is not supported in the following regions: China (Heyuan), China (Guangzhou), China (Hohhot), China (Ulanqab), UAE (Dubai), and Malaysia (Kuala Lumpur).
- You can configure up to 10 event notification rules in a region.
- Notifications are not sent for TS and M3U8 objects that are generated by ingesting streams over Real-Time Messaging Protocol (RTMP). For more information about RTMP-based stream ingest, see Overview.

For more information about the event notification feature, see Overview.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket

for which you want to configure event notification rules.

3. In the left-side navigation pane, choose **Basic Settings > Event Notification**.
4. In the Event Notification section, click **Configure**. On the page that appears, click **Create Rule**.
5. In the **Create Rule** panel, configure the rule parameters that are described in the following table.

Parameter	Description
Rule Name	Specify the name of the event notification rule. The name of each event notification rule that is created by using the same Alibaba Cloud account must be unique in the same region. The name of an event notification rule must start with a letter and can contain only letters, digits, and hyphens (-). The name cannot exceed 85 characters in length.
Events	Select one or more events that can trigger the event notification rule from the drop-down list. For example, if you want to receive a notification when a specific object is created or overwritten by copying an object, select CopyObject. You can configure an event notification rule for a specific object and specify multiple types of events that can trigger the rule. You can also configure multiple event notification rules for an object. When you configure multiple event notification rules, take note of the following items: ○ If the multiple event notification rules apply to the same object, the values of this parameter in these rules must be different. For example, if you select CopyObject for Events when you create an event notification rule for objects whose names contain the `images` prefix, CopyObject cannot be selected for Events when you create another event notification rule for objects whose names contain the same prefix. ○ If the multiple event notification rules apply to different objects, the values of this parameter in these rules can be the same. For example, if you select PutObject for Events when you create an event notification rule for objects whose names contain the `images` prefix and the `.png` suffix, you can select PutObject or DeleteObject for Events when you create another event notification rule for objects whose names contain the `log` prefix and the `.jpg` suffix. **Notice** If you do not specify the version ID when you delete an object from a versioned bucket, the DeleteObject or DeleteObjects event notification is not triggered. In this case, no version of the object is deleted. The current version of the object is converted into a previous version and a delete marker is added to the object. For more information about the object operations that correspond to the event types, see Events.

Parameter	Description
Resource Description	Specify the objects to which the event notification rule applies. ○ Select **Full Name** to apply the rule to an object whose name matches the specified name. ■ To create a rule that applies to an object named exampleobject.txt in the root directory of the bucket, enter *exampleobject.txt*. ■ To create a rule that applies to an object named myphoto.jpg in the destdir directory within the root directory of the bucket, enter *destdir/myphoto.jpg*. ○ Select **Prefix and Suffix** to apply the rule to objects whose names contain the specified prefix and suffix. ■ To create a rule that applies to all objects in the bucket, leave Prefix and Suffix empty. ■ To create a rule that applies to all objects in the examplefolder directory within the root directory of the bucket, set Prefix to *examplefolder/* and leave Suffix empty. ■ To create a rule that applies to all JPG objects in the bucket, leave Prefix empty and set Suffix to *.jpg*. ■ To create a rule that applies to all MP3 objects in the *examplefolder* directory within the root directory of the bucket, set Prefix to *examplefolder/* and Suffix to *.mp3*. To create a Resource Description entry, click **Add**. You can create up to five **Resource Description** entries.
Endpoint	Specify the endpoint to which notifications are sent. Valid values: **HTTP** and **Queue**. ○ **HTTP**: Enter the address of the HTTP endpoint to which notifications are sent. Example: `http://198.51.100.1:8080` . For more information about how to enable an HTTP endpoint, see Manage topics and HttpEndpoint. ○ **Queue**: Enter the name of an MNS queue. For more information about how to create a queue, see Create a queue. To create an endpoint, click **Add**. You can create up to five **endpoints**.

6. Click **OK**.

   After you configure the event notification rule, the rule takes effect after approximately 10 minutes.

2.5.10. Delete a bucket

If you no longer use a bucket, delete the bucket to stop unexpected charges.

> **Warning** Deleted buckets cannot be restored. Exercise caution when you perform this operation.

Prerequisites

- All objects in the bucket are deleted.
  - For more information about how to manually delete a small number of objects, see Delete objects.
  - To delete a large number of objects, we recommend that you configure lifecycle rules to batch delete the objects. For more information about how to delete a large number of objects, see Configure lifecycle rules.

  > Notice
  >
  > If versioning is enabled for the bucket that you want to delete, make sure that all versions of objects in the bucket are deleted. For more information about how to delete all versions of objects in a bucket, see Configure versioning.

- Parts that are uploaded by multipart upload or resumable upload tasks in the bucket are deleted. For more information about how to delete parts in a bucket, see Manage parts.
- All LiveChannels in the bucket are deleted. For more information about how to delete LiveChannels, see DeleteLiveChannel.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Basic Settings > Delete Bucket**.
4. Click **Delete Bucket**. In the message that appears, click **OK**.

2.6. Redundancy for fault tolerance

2.6.1. Configure CRR

Cross-region replication (CRR) provides automatic and asynchronous (near real-time) replication of objects across buckets in different Object Storage Service (OSS) regions. CRR can also synchronize operations such as create, overwrite, and delete operations performed on objects from a source bucket to a destination bucket.

Usage notes

- Billing
  - You are charged for the traffic that is generated when you use CRR to replicate objects in OSS. For more information about the billing methods, see Traffic fees.
  - Each time an object is synchronized, OSS accumulates the number of requests, and you are charged for the requests. For more information about the billing methods, see API fees.
  - If you enable transfer acceleration, you are charged for the feature. For more information about the billing methods, see Transfer acceleration fees.
- Replication time

In CRR, data is replicated asynchronously in near real time. The time that is required to replicate data from the source bucket to the destination bucket may range from a few minutes to a few hours. The replication time varies based on the data size.

Limits

- Limits on regions
  - CRR is unavailable in the South Korea (Seoul) and Thailand (Bangkok) region. For more information about the regions in which OSS data centers are located, see Regions and endpoints.
  - You must enable transfer acceleration when you perform CRR between the regions inside and outside the Chinese mainland.
  - CRR rules based on object tags can be configured only in the following scenarios:
    - The source region is China (Hangzhou), and the destination region is a region except for China (Hangzhou).
    - The source region is Australia (Sydney), and the destination region is a region outside the Chinese mainland and except for Australia (Sydney).
- Limits on operations
  - You can configure CRR between two unversioned buckets or two versioned buckets.
  - The versioning status of two buckets between which a CRR rule is configured cannot be changed.
  - If you configure a CRR rule for two buckets, an object replicated from the source bucket may overwrite an object that has the same name in the destination bucket.
  - You can configure CRR rules to synchronize data from a source bucket to multiple destination buckets. You can configure up to 100 CRR rules for a bucket. You can specify a bucket as a source bucket or a destination bucket. If you want to configure more than 100 CRR rules for a bucket, contact technical support.
  - Cold Archive objects in the source bucket cannot be synchronized to the destination bucket.
  - Appendable objects cannot be synchronized from a source bucket to a destination bucket whose storage class is Cold Archive.

Enable CRR

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**, and then click the name of the bucket for which you want to enable CRR.
3. In the left-side navigation pane, choose **Redundancy for Fault Tolerance > Cross-Region Replication**. In the **Cross-Region Replication** section, click **Configure**.
4. Click **Cross-Region Replication**. In the **Cross-Region Replication** panel, configure the parameters. The following table describes the parameters.

Parameter	Description
Source Region	Specify the region in which the current bucket is located.
Source Bucket	Specify the name of the current bucket.
Destination Region	Specify the region in which the destination bucket is located.

Parameter	Description
Destination Bucket	Specify the destination bucket to which you want to synchronize data.
Acceleration Type	Specify the acceleration type. Only **Transfer Acceleration** is supported. You can use transfer acceleration to accelerate data transfer when you replicate data across regions in the Chinese mainland and outside the Chinese mainland. If you enable transfer acceleration, you are charged for the use of this feature. For more information about the billing methods, see Transfer acceleration fees.
Applied To	Specify the source data that you want to synchronize. ○ **All Files in Source Bucket**: OSS synchronizes all objects from the source bucket to the destination bucket. ○ **Files with Specified Prefix**: OSS synchronizes the objects whose names contain the specified prefix from the source bucket to the destination bucket. You can specify up to 10 prefixes.
Object Tagging	Specify the tags of objects that you want to synchronize to the destination bucket. Objects that have the specified tags are synchronized to the destination bucket. Select **Configure Rules** and add tags in key-value pairs. You can add up to 10 tags. When you configure this parameter, make sure that the following conditions are met: ○ Tags are configured for objects. For more information, see Configure object tagging. ○ Versioning is enabled for the source bucket and the destination bucket. ○ The Operations parameter is set to **Add/Change**. ○ If the source region is China (Hangzhou), the destination region can be any region except China (Hangzhou). If the source region is Australia (Sydney), the destination region can be any region outside the Chinese mainland except Australia (Sydney).
Operations	Specify the synchronization policy. ○ **Add/Change**: OSS synchronizes the data changes including the create and overwrite operations on objects from the source bucket to the destination bucket. ○ **Add/Delete/Change**: OSS synchronizes all data changes including the create, overwrite, and delete operations on objects from the source bucket to the destination bucket. If you use the multipart upload method to upload an object to the source bucket, each uploaded part is synchronized to the destination bucket. The complete object that is obtained by calling the CompleteMultipartUpload operation is also synchronized to the destination bucket. For more information about how to configure CRR for objects in versioned buckets, see CRR in specific scenarios.

Parameter	Description
Replicate Historical Data	Specify whether to synchronize historical data in the source bucket before you enable CRR for the source bucket. ◦ **Yes**: OSS synchronizes historical data to the destination bucket. **Notice** When historical data is synchronized, objects in the source bucket may overwrite objects that have the same names in the destination bucket. To prevent data loss, we recommend that you enable versioning for the source and destination buckets. ◦ **No**: OSS synchronizes only objects that are uploaded or updated after the CRR rule takes effect to the destination bucket.
KMS-based Encryption	If KMS-based encryption is configured for the source objects or destination bucket, you must select **KMS-based Encryption** and configure the following parameters: ◦ **CMK ID**: specifies a customer master key (CMK) that is used to encrypt the destination object. If you want to use a CMK to encrypt objects, you must create a CMK in the same region as the destination bucket in the Key Management Service (KMS) console. For more information, see Create a CMK. ◦ **RAM Role Name**: specifies a RAM role that is authorized to perform KMS-based encryption on the destination object. ■ **New RAM Role**: A RAM role is created to perform KMS-based encryption on the destination object. The name of the RAM role is in the following format: `kms-replication-source bucket name-destination bucket name`. ■ **AliyunOSSRole**: The AliyunOSSRole role is used to perform KMS-based encryption on the destination object. If the AliyunOSSRole role does not exist, OSS automatically creates the AliyunOSSRole role when you select this option. **Note** ◦ You can use HeadObject to query the encryption status of the source object and use GetBucketEncryption to query the encryption status of the destination bucket. ◦ For more information about how to configure CRR for buckets for which server-side encryption is configured, see CRR in specific scenarios.

5. Click **OK**.
   - After you create a CRR rule, the rule cannot be edited or deleted.
   - After you configure a CRR rule, the synchronization task starts in 3 to 5 minutes. To view the synchronization progress, choose **Redundancy for Fault Tolerance > Cross-Region Replication** on the management page of the source bucket.

  - In CRR, data is replicated asynchronously. Depending on the amount of data, it can take a few minutes to several hours to replicate data to the destination bucket.

Disable CRR

You can click **Disable** to disable CRR.

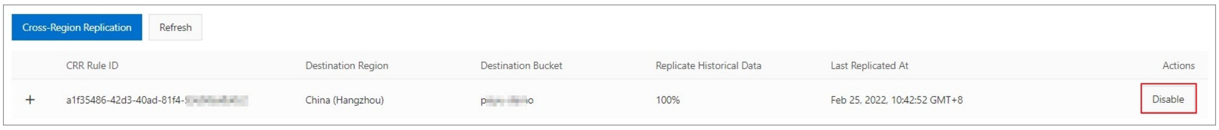

	CRR Rule ID	Destination Region	Destination Bucket	Replicate Historical Data	Last Replicated At	Actions
+	a1f35486-42d3-40ad-81f4-[redacted]	China (Hangzhou)	p[redacted]o	100%	Feb 25, 2022, 10:42:52 GMT+8	Disable

After you disable CRR, the replicated data is stored in the destination bucket. However, the incremental data in the source bucket is not replicated to the destination bucket.

2.6.2. Configure SRR

Same-region replication (SRR) allows you to replicate objects across buckets within the same region in an automatic and asynchronous (near real-time) manner. Operations such as the creation, overwriting, and deletion of objects can be synchronized from a source bucket to destination buckets.

Usage notes

- Billing

  After SRR is enabled, you are not charged for the traffic that is generated when you use SRR to replicate objects from the source bucket to the destination bucket in Object Storage Service (OSS). Each time an object is synchronized, OSS accumulates the number of requests. However, you are not charged for the requests.

- Replication time

  In SRR, data is replicated asynchronously in near real time. The time required to replicate data from the source bucket to the destination bucket may be a few minutes to several hours. The replication time varies based on the data size.

Usage notes

- You can configure CRR between two unversioned buckets or two versioned buckets.
- The versioning status of two buckets between which a CRR rule is configured cannot be changed.
- If you configure a CRR rule for two buckets, an object replicated from the source bucket may overwrite an object that has the same name in the destination bucket.
- Cold Archive objects in the source bucket cannot be synchronized to the destination bucket.
- Appendable objects cannot be synchronized from a source bucket to a destination bucket whose storage class is Cold Archive.

Enable SRR

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the page that appears, click the name of the bucket for which you want to enable SRR.
3. In the left-side navigation pane, choose **Redundancy for Fault Tolerance > Same-Region Replication**.
4. In the **Same-Region Replication** section, click **Configure**.

5. Click **Same-Region Replication**.
6. In the **Same-Region Replication** panel, configure the parameters described in the following table.

Parameter	Description
Source Region	The region in which the current bucket is located.
Source Bucket	The name of the current bucket.
Destination Bucket	Select the destination bucket to which you want to synchronize data.
Applied To	Select the source data that you want to synchronize. ◦ **All Files in Source Bucket**: OSS synchronizes all objects from the source bucket to the destination bucket. ◦ **Files with Specified Prefix**: OSS synchronizes the objects whose names contain a specified prefix from the source bucket to the destination bucket. You can specify up to 10 prefixes.
Object Tagging	The tags of objects that you want to synchronize to the destination bucket. Objects that have the specified tags are synchronized to the destination bucket. Select **Configure Rules** and add tags in key-value pairs. You can add up to 10 tags. To configure this parameter, make sure that the following conditions are met: ◦ Tags are configured for objects. For more information, see Configure object tagging. ◦ Versioning is enabled for the source bucket and the destination bucket. ◦ The Operations parameter is set to **Add/Change**.
Operations	Select the operations to synchronize. ◦ **Add/Change**: OSS synchronizes only the added or changed data from the source bucket to the destination bucket. ◦ **Add/Delete/Change**: OSS synchronizes all data changes including the create, overwrite, and delete operations on objects from the source bucket to the destination bucket.
Replicate Historical Data	Specifies whether to synchronize historical data in the source bucket before you enable SRR for the source bucket. ◦ **Yes**: OSS synchronizes historical data to the destination bucket. **Notice** When historical data is synchronized, objects in the source bucket may overwrite objects that have the same names in the destination bucket. To avoid data loss, we recommend that you enable versioning for the source and destination buckets. ◦ **No**: OSS synchronizes only objects that are uploaded or updated after the SRR rule takes effect to the destination bucket.

Parameter	Description
KMS-based Encryption	If KMS-based encryption is configured for the source objects or destination bucket, you must select **KMS-based Encryption** and configure the following parameters: ○ **CMK ID**: The customer master key (CMK) that is used to encrypt the destination object. If you want to use a CMK to encrypt objects, you must create a CMK in the same region as the destination bucket in the Key Management Service (KMS) console. For more information, see Create a CMK. ○ **RAM Role Name**: The RAM role that is authorized to perform KMS-based encryption on the destination object. ■ **New RAM Role**: A RAM role is created to perform KMS-based encryption on the destination object. The name of the RAM role is in the following format: `kms-replication-source bucket name-destination bucket name`. ■ **AliyunOSSRole**: The AliyunOSSRole role is used to perform KMS-based encryption on the destination object. If the AliyunOSSRole role does not exist, OSS automatically creates the AliyunOSSRole role when you select this option. **Note** You can use HeadObject to query the encryption status of the source object and use GetBucketEncryption to query the encryption state of the destination bucket.

7. Click **OK**.
   - An SRR rule cannot be edited or deleted after it is created.
   - The synchronization starts immediately after an SRR rule is configured. You can view the synchronization progress on the **Same-Region Replication** page.
   - It can take several minutes to several hours for the data to be synchronized to the destination bucket based on the amount of data.

Disable SRR

You can click **Disable** to disable SRR.

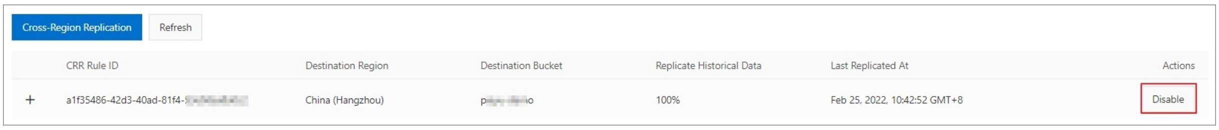

	CRR Rule ID	Destination Region	Destination Bucket	Replicate Historical Data	Last Replicated At	Actions
+	a1f35486-42d3-40ad-81f4-[redacted]	China (Hangzhou)	p[redacted]o	100%	Feb 25, 2022, 10:42:52 GMT+8	Disable

After you disable SRR, the replicated data is stored in the destination bucket. However, the incremental data in the source bucket is not synchronized to the destination bucket any more.

2.6.3. Configure versioning

Object Storage Service (OSS) allows you to configure versioning to protect data at the bucket level. After you enable versioning for a bucket, data that is overwritten or deleted in the bucket is saved as a previous version. You can use versioning to restore a previous version of an object that is accidentally overwritten or deleted.

Context

If you enable versioning for a bucket, you are charged for the storage of all versions of objects in the bucket. If you download or restore a previous version of an object, you are charged for the requests and traffic. To prevent unnecessary storage fees, we recommend that you delete the previous versions of objects that you no longer need at your earliest opportunity. For more information, see Billing overview.

Enable versioning

When versioning is enabled for a bucket, OSS specifies a unique ID for each version of an object stored in the bucket.

- Enable versioning when you create a bucket
    i. Log on to the OSS console.
    ii. In the left-side navigation pane, click **Buckets**. On the Buckets page, click **Create Bucket**.
    iii. In the **Create Bucket** panel, configure the parameters.

       Set **Versioning** to **Activate**. For more information about how to configure other parameters, see Create buckets.
    iv. Click **OK**.
- Enable versioning for an existing bucket
    i. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to enable versioning.
    ii. In the left-side navigation pane, choose **Redundancy for Fault Tolerance > Versioning**.
    iii. Click **Configure**. Set Versioning to **Enabled**.
    iv. Click **Save**.

After versioning is enabled for a bucket, you can view all versions of objects in the bucket on the **Files** page. If you want to view only the current versions of objects, set **Display Previous Versions** to **Hide**. Hiding previous versions of objects does not increase the response speed when you list objects. If the response speed is slow when you list objects, see Low response speed to troubleshoot and fix the problem.

Restore a previous version

To restore a specified previous version of an object as the current version, perform the following steps:

1. In the left-side navigation pane, click **Files**.
2. Restore a specified previous version of an object as the current version.

   > **Notice** You can restore only one previous version of an object at a time. The previous version that you want to restore cannot be a delete marker.

   - Restore the previous version of an object

Click **Restore** in the Actions column that corresponds to the previous version that you want to restore.

- Restore the previous versions of multiple objects

  Select the previous versions that you want to restore and then choose **Batch Operation > Restore**.

Download a specified version of an object

To download a specified version of an object, perform the following steps:

1. In the left-side navigation pane, click **Files**.
2. Click the version that you want to download. In the panel that appears, click **Download** on the right side of **Object URL**.
3. Select the location where you want to store the downloaded version and then click **Save**.

Delete a previous version of an object

To minimize storage costs, we recommend that you delete the previous versions of objects that you no longer need at your earliest opportunity.

> Warning
>
> - You cannot restore a previous version of an object after it is deleted. Proceed with caution.
> -

1. In the left-side navigation pane, click **Files**.
2. Click **Permanently Delete** in the Actions column that corresponds to the previous version you want to delete.

   To delete multiple previous versions of an object, select the previous versions that you want to delete and choose **Batch Operation > Permanently Delete**.
3. Click **OK**.

You can also configure lifecycle rules to allow OSS to periodically delete previous versions. For more information, see Configure lifecycle rules.

Suspend versioning

You can suspend versioning for a versioned bucket to stop OSS from generating new versions for objects. If a new version is generated for an object in a versioning-suspended bucket, OSS sets the ID of the new version to null and retains the previous versions of the object.

To suspend versioning for a bucket, perform the following steps:

1. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to suspend versioning.
2. In the left-side navigation pane, choose **Redundancy for Fault Tolerance > Versioning**.
3. Click **Configure** and set the versioning state to **Suspended**.
4. Click **Save**.

2.7. Manage logs

2.7.1. Configure logging

A large number of logs are generated when your Object Storage Service (OSS) buckets are accessed. After you enable and configure logging for a bucket, OSS generates logs every hour in accordance with predefined naming conventions and then stores the logs as objects in a specified bucket. You can use Alibaba Cloud Log Service or build a Spark cluster to analyze the logs.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure logging.
3. In the left-side navigation pane, choose **Logging** > **Logging**.
4. Click **Configure**. Set Destination Bucket and Log Prefix.
    - **Destination Bucket**: Select a bucket name from the Destination Bucket drop-down list. You can only select a bucket that is located in the same region as the bucket for which logging is enabled within the same Alibaba Cloud account.
    - **Log Prefix**: Enter the path and prefix of the logs. If you specify this parameter, log objects are stored in a specified directory of the destination bucket. If you do not specify this parameter, log objects are stored in the root directory of the destination bucket. For example, if you enter *log/* in the Log Prefix field, the log objects are stored in the *log/* directory.
5. Click **Save**.

2.7.2. Query real-time logs

When you access Object Storage Service (OSS) resources, a large number of logs are generated. OSS uses Log Service to help you query and collect statistics for OSS access logs and audit access to OSS in the OSS console, track exception events, and troubleshoot problems. This helps you improve work efficiency and make informed decisions.

Prerequisites

- Log Service is activated.

  If you have not activated Log Service, go to the Log Service product page.

- Log Service is authorized to access OSS.

  If you have not authorized Log Service to access OSS, visit Cloud Resource Access Authorization and follow the instructions to complete the authorization.

Billing

When you use the real-time log feature, fees are generated if the following conditions occur:

- Real-time log query allows you to query logs over the last seven days free of charge. If the log retention time that you set is longer than seven days, you are charged for the excess days.
- Log Service allows you to store 900 GB of logs (equivalent to 900 million 1-KB log entries) per day free of charge. You are charged for the excess logs.
- You are charged for traffic consumed when you read data from or write data to Log Service over the Internet.

For more information about the billing standards, see Pay-as-you-go.

Enable real-time log query

You can use one of the following methods to enable real-time log query:

- i. Log on to the OSS console.

  ii. On the **Overview** page, click **Create Bucket** on the right side.

  iii. In the **Create Bucket** dialog box, set **Real-time Log Query** to **Enable**. For more information about other parameters, see Create buckets.

  iv. Click **OK**.

Method 1: Enable real-time log query when you create a bucket

- i. Log on to the OSS console.

  ii. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to enable real-time log query.

  iii. In the left-side navigation pane, choose **Logging > Real-time Log Query**.

  iv. Click **Activate Now**.

Method 2: Enable real-time log query for an existing bucket

Real-time log query allows you to query logs over the last seven days free of charge. You can click **Config Log Retention Time** in the upper-right corner to modify the retention time of logs.

Query real-time logs

You can use one of the following methods to query real-time logs:

- i. Log on to the OSS console.

  ii. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket whose logs you want to query.

  iii. In the left-side navigation pane, choose **Logging > Real-time Log Query**.

  iv. Click **Original Log** to analyze logs.

You can specify the time range and query statement in real-time log queries. For example, you can analyze the distribution of a specified field such as an API operation within a specified time range. You can also filter the query results by conditions to view required access records.

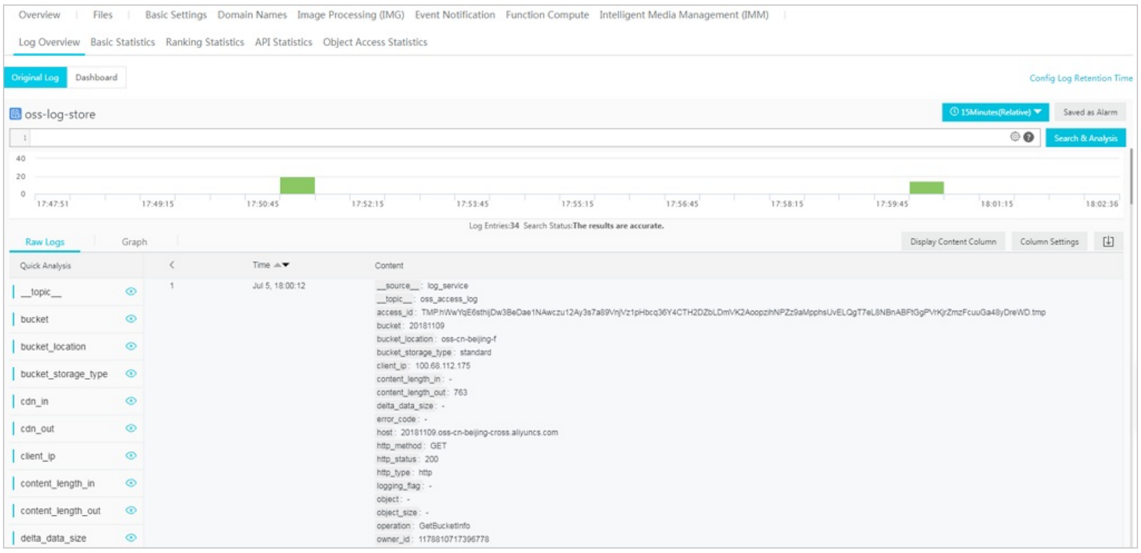

Log Entries:34 Search Status:**The results are accurate.**

```
1    Jul 5, 18:00:12    __source__ : log_service
                        __topic__ : oss_access_log
                        access_id : TMP.hWwYqE6sthjDw3BeDae1NAwczu12Ay3s7a89VnjVz1pHbcq36Y4CTH2DZbLDmVK2AoopzihNPZz9aMpphsUvELQgT7eL8NBnABFtGgPVrKjrZmzFcuuGa48yDreWD.tmp
                        bucket : 20181109
                        bucket_location : oss-cn-beijing-f
                        bucket_storage_type : standard
                        client_ip : 100.68.112.175
                        content_length_in : -
                        content_length_out : 763
                        delta_data_size : -
                        error_code : -
                        host : 20181109.oss-cn-beijing-cross.aliyuncs.com
                        http_method : GET
                        http_status : 200
                        http_type : http
                        logging_flag : -
                        object : -
                        object_size : -
                        operation : GetBucketInfo
                        owner_id : 1178810717396778
```

Method 1: Query real-time logs on the Original Log page

- i. Log on to the OSS console.

  ii. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket whose logs you want to query.

  iii. In the left-side navigation pane, choose **Logging > Real-time Log Query**.

  iv. Click **Dashboard** to analyze logs.

  Dashboard allows you to view the following types of reports:

  - **Access Center**: displays the overall operating status including the PV, UV, traffic, and access distribution over the Internet.
  - **Audit Center**: displays statistics on object operations including read, write, and delete operations.
  - **Operation Center**: displays statistics on access logs including the number of requests and distribution of failed operations.
  - **Performance Center**: displays statistics on performance including the performance of downloads and uploads over the Internet, the performance of transmission over different networks or with different object sizes, and the list of differences between stored and downloaded objects.

Method 2: Query real-time logs on the Dashboard page

- Log on to the Log Service console to query real-time OSS logs. For more information, see OSS access logs.

Method 3: Query real-time logs in the Log Service console

Disable real-time log query

If you no longer need to retain log data, perform the following steps to disable real-time log query:

1. Log on to the OSS console.

2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket whose logs you want to disable.
3. Choose **Logging > Real-time Log Query**.
4. Click **Disable Real-time Log Query** in the upper-right corner.

   > **Notice** When you enable real-time log query, Log Service projects are automatically created. However, the projects are not automatically deleted when you disable real-time log query. Therefore, after real-time log query is disabled, you must log on to the Log Service console to delete the projects that are automatically created to avoid unexpected charges. For more information, see Delete a project.

References

- For more information about log query, see Real-time log query.
- For more information about OSS logging, see Configure logging.

2.8. Data Processing

2.8.1. Configure image styles

You can encapsulate multiple image processing (IMG) parameters in a style and perform complex IMG operations by using the style.

Context

Up to 50 styles can be created for a bucket. These styles can be used only for image objects in the bucket. If you want to create more than 50 styles, contact technical support.

Create a style

1.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the bucket that you want to configure.
3. In the left-side navigation pane, choose **Data Processing > Image Processing (IMG)**. Then, click **Create Rule**.
4. In the **Create Rule** pane, configure the style.

   You can use **Basic Edit** or **Advanced Edit** to create a style:

   - **Basic Edit**: You can use the IMG parameters listed by using the graphical user interface (GUI) to choose the IMG methods. For example, resize an image, add a watermark, and modify the image format.
   - **Advanced Edit**: You can use the API code to edit the IMG methods. The format is: `image/action1,parame_value1/action2,parame_value2/...` . For more information about supported IMG parameters, see the "Parameters" section of the Overview topic.

     Example: `image/resize,p_63/quality,q_90` indicates that the image is scaled down to 63% of the source image, and then the relative quality of the image is set to 90%.

> Note If you want to add image and text watermarks to images at the same time by using a style, use **Advanced Edit** to create the style.

5. Click **OK**.

Apply styles

After the style is created, you can use the style in the bucket to process your image objects.

1. On the buckets page, click **Files**.
2. Click the name of the image.
3. In the **View Details** panel, select an image style from the **Image Style** drop-down list.

   You can view the processed image in the **View Details** panel. Right-click the image and click **Save As** to save the image to your local computer.

   You can also add styles to the IMG URLs and SDKs. For more information, see Usage notes.

Import the styles of the source bucket to the destination bucket

You can export styles that are created in the source bucket and import the styles to the destination bucket. This way, you can apply styles in the destination bucket to process image objects.

1. Export styles in the source bucket.
   i. On the Overview page of the source bucket, choose **Data Processing > Image Processing (IMG)** in the left-side navigation pane.
   ii. Click **Export**.
   iii. In the **Save As** dialog box, select the path to which you want to save the style, and click **Save**.
2. Import styles to the destination bucket.
   i. On the Overview page of the destination bucket, choose **Data Processing > Image Processing (IMG)** in the left-side navigation pane.
   ii. Click **Import**.
   iii. In the **Open** dialog box, select the exported style file and click **Open**.

      After the style is imported to the destination bucket, you can use the style to process images stored in the bucket.

Simplify IMG URLs that have style parameters

IMG URLs that have styles contain file access URLs, style parameters, and style names. Example: https://image-demo-oss-zhangjiakou.oss-cn-zhangjiakou.aliyuncs.com/example.jpg? x-oss-process=style/small. You can replace the `?x-oss-process=style/` field by using custom delimiters to simplify the IMG URL. For example, set the custom delimiter to an exclamation point (!). IMG URLs can be replaced by https://image-demo-oss-zhangjiakou.oss-cn-zhangjiakou.aliyuncs.com/example.jpg!small.

1. On the Overview page of the bucket, choose **Data Processing > Image Processing (IMG)** in the left-side navigation pane.
2. Click **Access Settings**.
3. On the **Access Settings** pane, select **Delimiters**.

Only hyphens (-), underscores (_), forward slashes (/), and exclamation points (?) can be used as delimiters.

4. Click **OK**.

   You can also bind a custom domain name to the bucket to further simplify the IMG URL. For example, if the sample bucket is bound to the `example.com` custom domain name , the sample URL can be replaced with `https://example.com/example.jpg! small` . After you bind a custom domain name, you can also preview the effect of IMG online. For more information, see Map custom domain names.

References

- For more information about how to use IMG parameters to process images, see IMG implementation modes.
- For more information about how to save processed images in OSS, see Save processed images.

3.Upload, download, and manage objects

3.1. Overview

Objects are basic units for user operations in OSS. The maximum size of an object is 48.8 TB. The number of objects stored in a bucket is not limited.

After you create a bucket in a region, the objects uploaded to the bucket are retained in this region, unless you manually migrate the objects to another region. You can access the objects anywhere if you have access permissions.

You must have write permissions on a bucket to upload objects to the bucket. Uploaded objects are displayed as files or folders in the OSS console. This topic describes how to create, manage, and delete objects and folders in the OSS console.

3.2. Upload objects

You can upload objects up to 5 GB in size in the Object Storage Service (OSS) console.

Prerequisites

A bucket is created. For more information, see Create buckets.

Context

To upload objects larger than 5 GB in size, we recommend that you use Multipart upload of OSS SDKs or OSS APIs, graphical management tool ossbrowser, or command-line tool ossutil.

In Alibaba Finance Cloud, you cannot upload objects by using the OSS console because OSS does not have a region connected to the Internet. To upload objects, you must use OSS SDKs, ossutil, or ossbrowser.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket to which you want to upload objects.
3. In the left-side navigation pane, click the **Files** tab. On the page that appears, click **Upload**.
4. In the **Upload** panel, configure the parameters described in the following table.
    i. The following table describes the basic settings.

Parameter	Description

Parameter	Description
Upload To	Set the path in which to store an object after the object is uploaded to the bucket. - **Current**: Objects are uploaded to the current directory. - **Specified**: Objects are uploaded to the specified directory. You must enter the directory name. If the directory whose name you entered does not exist, OSS automatically creates the directory and uploads the object to the directory. The directory must meet the following naming conventions: - The name can contain only UTF-8 characters. The name must be 1 to 254 characters in length. - The name cannot start with a forward slash (/) or backslash (\). - The name cannot contain consecutive forward slashes (/). - The name of the directory cannot be `..`.
File ACL	Set the access control list (ACL) for the object. - **Inherited from Bucket**: The ACL of the object is the same as that of the bucket. - **Private**: Only the object owner or authorized users can read and write the objects to upload. Other users, including anonymous users, cannot access the objects without authorization. We recommend that you set the File ACL parameter to this value. - **Public Read**: Only the owner or authorized users of this bucket can write the objects to upload. Other users, including anonymous users, can only read the objects. If you set the File ACL parameter to this value, the objects may be unexpectedly accessed, which results in out-of-control costs. - **Public Read/Write**: All users, including anonymous users, can read and write the objects to upload. If you set the File ACL parameter to public read/write, the objects may be unexpectedly accessed, which results in out-of-control costs. If a user uploads prohibited data or information, your legitimate interests and rights may be infringed. Therefore, we recommend that you do not set the File ACL parameter to public read/write except in special cases. For more information about object ACLs, see Object ACL.
Upload Acceleration	After transfer acceleration is enabled for the bucket that contains the object, you can turn on **Upload Acceleration** if you want to accelerate the upload of the object. For more information about transfer acceleration, see Transfer acceleration.

Parameter	Description
Files to Upload	Select the files or directories that you want to upload. You can click **Select Files** to select a local file or click **Select Folders** to select a directory. You can also drag the required local file or directory to the Files to Upload section. If you select an unnecessary object, click **Remove** in the Actions column that corresponds to the object to remove the object. Notice - When you upload an object that has the same name as an existing object in OSS to an unversioned bucket, the existing object is overwritten. - When you upload an object that has the same name as an existing object in OSS to a versioned bucket, the existing object becomes a previous version, and the uploaded object becomes the latest version.

ii. (Optional) Configure advanced settings such as Storage Class and Encryption Method.

Parameter	Description
Storage Class	Set the storage class of the object. - **Inherited from Bucket**: The storage class of the object is the same as that of the bucket. - **Standard**: Standard is suitable for objects that are frequently accessed. - **IA**: Infrequent Access (IA) is suitable for objects that are less frequently accessed. On average, objects that are accessed less than once to twice a month fall into this category. IA objects have a minimum storage duration of 30 days. You are charged for data retrieval when you access these objects. - **Archive**: Archive is suitable for objects that are infrequently accessed. Archive objects have a minimum storage duration of 60 days. Before you can access an object of the Archive storage class, you must restore the object. The restoration takes about one minute, and data retrieval fees are incurred during the restoration process. - **Cold Archive**: Cold Archive is suitable for long-term storage of backup objects and raw data. Cold Archive objects have a minimum storage duration of 180 days. Before you can access an object of the Cold Archive storage class, you must restore the object. The amount of time required to restore a Cold Archive object depends on the data size and the restore mode. You are charged for the data retrieval when you restore a Cold Archive object. For more information, see Overview.

Parameter	Description
Encryption Method	Configure server-end encryption method for an object. ■ **Inherited from Bucket**: The encryption method of the object is the same as that of the bucket. ■ **OSS-Managed**: Keys managed by OSS are used to encrypt objects in the bucket. OSS uses data keys to encrypt objects. In addition, OSS uses regularly rotated master keys to encrypt data keys. ■ **KMS**: The default CMK stored in Key Management Service (KMS) or the specified CMK ID is used to encrypt and decrypt data. Descriptions of **CMK**: ■ **alias/acs/oss**: The default customer master key (CMK) stored in KMS is used to encrypt different objects and decrypt the objects when the objects are downloaded. ■ CMK ID: The keys generated by a specified CMK are used to encrypt different objects and the specified CMK ID is recorded in the metadata of the encrypted object. Objects are decrypted when they are downloaded by users who are granted decryption permissions. Before you specify a CMK ID, you must create a normal key or an external key in the same region as the bucket in the KMS console. ■ **Encryption algorithm**:Only AES-256 is supported. For more information about object ACLs, see Object ACL.
User Metadata	Add the descriptive information for the object. You can add multiple pieces of user metadata as custom headers. However, the total size of the user metadata cannot exceed 8 KB. When you add user metadata, you must prefix parameters with `x-oss-meta-` and specify a value such as **x-oss-meta-location:hangzhou** for the parameters.

iii. Click **Upload**.
You can view the upload progress of objects on the **Task List** tab.

3.3. Create directories

If you want to organize the objects that you upload to a bucket, you can create directories. This topic describes how to create directories in the Object Storage Service (OSS) console.

Prerequisites

A bucket is created. For more information, see Create buckets.

Context

When the hierarchical namespace feature is disabled for a bucket, OSS uses a flat structure instead of a hierarchical structure for objects. All elements are stored as objects in buckets. To help organize objects and simplify management, the OSS console displays objects whose names end with a forward slash (/) as directories. The objects can be uploaded and downloaded. You can use directories in the OSS console in the same manner as you use directories in Windows.

When the hierarchical namespace feature is enabled for a bucket, you can update a parent directory to perform multiple directory-level tasks at the same time. Traditionally, OSS uses a flat namespace to store objects in buckets and uses objects whose names end with a forward slash (/) to simulate directories. Compared with this method, the hierarchical namespace feature greatly improves the performance of operations that are related to directories. For more information about the hierarchical namespace feature, see Hierarchical namespace.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which you want to create a directory.
3. In the left-side navigation pane, choose **Files** > **Files**.
4. On the **Files** page, click **Create Folder**.
5. In the **Create Folder** panel, enter a directory name in the **Folder Name** field.

   The directory name must meet the following conventions:

   - The directory name must be UTF-8 characters and cannot contain emoticons.
   - Forward slashes (`/`) are used in a directory name to indicate subdirectories. You can specify a directory name that contains multiple forward slashes (/) to create subdirectories in the directory. The directory name cannot start with a forward slash (`/`) or a backslash (`\`). The directory name cannot contain consecutive forward slashes (`/`).
   - The name of a subdirectory cannot be two consecutive periods (`..`).
   - The directory name must be 1 to 254 characters in length.
6. Click **OK**.

3.4. Configure ACL for objects

OSS allows you to configure ACL for objects to control access to them. This topic describes how to configure ACL for an object through the OSS console.

Prerequisites

- A bucket has been created. For more information, see Create a bucket.
- Objects have been uploaded to the bucket. For more information, see Upload objects.

Context

You can configure ACL for an object when or after you upload the object. If you do not specify ACL for an object, default ACL **Inherited from Bucket** applies.

- **Inherited from Bucket**: The ACL for object is the same as that for the bucket.
- **Private**: Only the bucket owner or authorized users can read from and write to the objects in the bucket. Other users, including anonymous users, cannot access objects in the bucket.

  > ? **Note** You can configure and send the object URL to share the object to other visitors. For more information, see Add signatures to URLs.

- **Public Read**: Only the bucket owner or authorized users can read from and write to objects. Other users, including anonymous users, can only read from objects in the bucket.

- **Public Read/Write**: All users, including anonymous users, can perform read and write operations on objects in the bucket. Fees incurred by these operations are paid by the bucket owner. Use this ACL policy only when necessary.

For more information about object ACL, see Object ACL.

Procedure

1.
2.
3. Click the **Files** tab.
4. Click the name of an object. In the **View Details** dialog box that appears, Click **Set ACL**.

   Alternatively, move the pointer over **More** in the Actions column corresponding to the object and choose **Set ACL** from the shortcut menu.
5. In the **Set ACL** dialog box that appears, set ACL for the object.
6. Click **OK**.

3.5. Modify the storage class of objects

Object Storage Service (OSS) provides the following storage classes: Standard, Infrequent Access (IA), Archive, and Cold Archive. Each storage class differs in the access frequency and billing. You can select a suitable storage class based on your business requirements. This topic describes how to modify the storage class of an object in the OSS console.

Usage notes

- When you modify the storage class of an object in the OSS console, the size of the object cannot exceed 1 GB. We recommend that you use ossutil to modify the storage classes of objects whose size are larger than 1 GB.
- You can modify the storage class of an Archive or Cold Archive object only after the object is restored. For more information, see Restore objects.
- When you modify the storage class of an object, the original object is replaced with a new object that has the selected storage class. If you change the storage class of an object to IA, Archive, or Cold Archive, you are charged storage fees based on the object size and storage duration when you access the IA object, or storage fees based on the object size and storage duration and data retrieval fees when you access the Archive or Cold Archive object. If the object size is smaller than 64 KB and the object is stored for a period of time that is shorter than the minimum storage duration, the minimum billable size 64 KB and the minimum storage duration are used for billing.
-

Procedure

1.
2.
3. In the left-side navigation pane, choose **Files** > **Files**.
4. On the **Files** page, move the pointer over **More** in the Actions column corresponding to the object

for which you want to modify the storage class and select **Modify Storage Class** from the drop-down list.

5. Select the storage class to which you want to convert, and click **OK**.

   We recommend that you keep **Retain User Metadata** enabled to retain the user metadata of the object after you modify the storage class.

3.6. Configure bucket policies to authorize other users to access OSS resources

You can configure bucket policies to grant permissions to other users to access specified Object Storage Service (OSS) resources.

Context

- The owner of a bucket can configure bucket policies for the bucket in the OSS console by using the GUI or by specifying policy syntax. Before you configure bucket policies by specifying policy syntax, you must understand the Action, Resource, and Condition fields in bucket policies. For more information, see Overview.
- If you select Anonymous Accounts (*) for the Accounts parameter and do not configure the Conditions parameter when you configure a bucket policy, the bucket policy applies to all users except for the bucket owner. If you select Anonymous Accounts (*) for the Accounts parameter and configure the Conditions parameter when you configure a bucket policy, the bucket policy applies to all users, including the bucket owner.
- You can configure multiple bucket policies for a bucket. The total size of the policies cannot exceed 16 KB.

Method 1: Configure bucket policies by using the GUI

1.
2.
3. In the left-side navigation pane, choose **Files > Files**, and click **Authorize**.

   You can also choose **Access Control > Bucket Policy** in the left-side navigation pane, and click **Configure** in the Bucket Policy section.

4. On the **GUI** tab, click **Authorize**.
5. In the **Authorize** panel, configure the parameters and click **OK**. The following table describes the parameters.

Parameter	Description

Parameter	Description
Applied To	Select the resources on which you want to grant other users the access permissions. ◦ **Whole Bucket**: The bucket policy applies to all resources in the bucket. ◦ **Specified Resource**: The bucket policy applies only to specified resources in the bucket. You can configure multiple bucket policies for specific resources in a bucket. ▪ Configure a bucket policy for a directory To configure a bucket policy to grant users the permissions to access all subdirectories and objects in a directory, add an asterisk (*) after the name of the directory. For example, to grant users the permissions to access all subdirectories and objects in a directory named abc, enter *abc/**. ▪ Configure a bucket policy for a specific object To configure a bucket policy to grant users the permissions to access a specific object, enter the full path of the object that excludes the bucket name. For example, to grant users the permissions to access an object named myphoto.png in the abc directory, enter *abc/myphoto.png*.

Parameter	Description
Accounts	Select the type of accounts to which you want to grant the permissions. ○ **Anonymous Accounts (*)**: Select this option if you want to grant all users the permissions to access the specified resources. ○ **RAM Users**: Select this option if you want to grant the RAM users of the current Alibaba Cloud account the permissions to access the specified resources. You can select individual RAM users from the drop-down list. If you want to grant the permissions to multiple RAM users, we recommend that you enter the keyword of the RAM usernames in the search box to perform fuzzy match. **Notice** If you select this option, you must log on to the OSS console by using an Alibaba Cloud account or a RAM user that has the management permissions on the bucket and the ListUsers permission in the RAM console. If you do not use an Alibaba Cloud account or a RAM user that has the required permissions, you cannot view the RAM user list of the current Alibaba Cloud account. For more information about how to grant the ListUsers permission to a RAM user, see Grant permissions to a RAM user. ○ **Other Accounts**: Select this option if you want to grant other Alibaba Cloud accounts, RAM users, or temporary users generated by Security Token Service (STS) the permissions to access the specified resources. ■ To grant other Alibaba Cloud accounts or RAM users the permissions to access the specified resources, enter the UIDs of the Alibaba Cloud accounts or RAM users. ■ To grant temporary users generated by STS the permissions to access the specified resources, enter the user and role information in the following format: `arn:sts::{RoleOwnerUid}:assumed-role/{RoleName}/{RoleSessionName}`. For example, the role used to generate a temporary user is testrole, the UID of the Alibaba Cloud account that assumes the role is 12345, and the RoleSessionName that is specified when the temporary user is generated is testsession. In this case, enter `arn:sts::12345:assumed-role/testrole/testsession`. To grant all temporary users the permissions to access the specified resources, use asterisks (*) as wildcard characters. For example, enter `arn:sts::*:*/*/*`. For more information about how to generate a temporary user, see Use a temporary credential provided by STS to access OSS. **Notice** If you grant a temporary user generated by STS the permissions to access your OSS resources, the temporary user cannot use the OSS console to access your OSS resources. However, the user can use ossutil, OSS API operations, or OSS SDKs to access your OSS resources.
	You can use the following methods to specify authorized operations: **Basic Settings** and **Advanced Settings**.

Parameter	Description
Authorized Operation	○ Basic Settings If you select this option, you can configure the following permissions based on your requirements. You can move the pointer over the ? icon on the right side of each permission to view the actions that correspond to the permission. ■ **Read-Only (excluding ListObject)**: allows authorized users to view and download the specified resources. ■ **Read-Only (including ListObject)**: allows authorized users to view, list, and download the specified resources. ■ **Read/Write**: allows authorized users to read data from and write data to the specified resources. ■ **Any Operation**: allows authorized users to perform all operations on the specified resources. ■ **None**: forbids authorized users from performing operations on the specified resources. Notice ■ ■ If multiple bucket policies are configured for a user, the user has all the permissions configured in the policies. However, the policy in which the Authorized Operation parameter is set to None takes precedence. For example, if you configure a first policy to grant the Read-Only permission to a user and configure a second policy to grant the Read/Write permission to the same user, the permission of the user is Read/Write. If you configure a third policy to grant the None permission to the user, the permission of the user is None. ■ The authorization effect for Read-Only (excluding ListObject), Read-Only (including ListObject), Read/Write, and Any Operation is Allow, and the authorization effect for None is Deny. ○ Advanced Settings If you select this option, you must configure the following parameters: ■ **Effect**: Select Allow or Deny. ■ **Action**: Specify the action that you want to allow or deny. You can specify an action that is supported by OSS. For more information about the actions that are supported by OSS, see Overview.

Parameter	Description
Conditions	Optional. You can configure this parameter in both Basic Settings and Advanced Settings to specify the conditions that users must meet before the users can access OSS resources. ○ **Access Method**: By default, authorized users can access OSS resources over both HTTP and HTTPS. If you want the authorized users to access the specified resources in the bucket over HTTPS, select **HTTPS**. If you want the authorized users to access the specified resources in the bucket over HTTP, select **HTTP**. Compared with HTTP, HTTPS is more secure. If you want to force all requests to access resources in the bucket by using one protocol, such as HTTPS, you must specify the syntax of the bucket policy. For more information, see How do I configure an HTTPS request and an SSL certificate? ○ **IP =**: Specify the IP addresses or CIDR blocks that can be used to access OSS resources. Separate multiple IP addresses with commas (,). ○ **IP ≠**: Specify the IP addresses or CIDR blocks that cannot be used to access OSS resources. Separate multiple IP addresses with commas (,). ○ **VPC**: Select the ID of the Apsara Stack VPC that belongs to the current Alibaba Cloud account from the drop-down list. You can also enter the ID of the VPC created by using the current account or another account in the field below. For more information about how to create a VPC, see Create a VPC.

6. Click **OK**.

Method 2: Configure bucket policies by specifying policy syntax

1.
2.
3. In the left-side navigation pane, choose **Files > Files**, and click **Authorize**.
4. On the **Syntax** tab, click **Edit**.

   You can specify policy syntax based on your business requirements for fine-grained access control. The following sample code provides examples on how the resource owner whose UID is `174649585760xxxx` configures the bucket policies in various scenarios:

   - Example 1: Allow anonymous users to list all objects in a bucket named examplebucket.

```
{
    "Statement": [
        {
            "Action": [
                "oss:ListObjects",
                "oss:ListObjectVersions"
            ],
            "Effect": "Allow",
            "Principal": [
                "*"
            ],
            "Resource": [
                "acs:oss:*:174649585760xxxx:examplebucket"
            ]
        },
    ],
    "Version": "1"
}
```

- Example 2: Forbid anonymous users whose IP addresses are not in the CIDR block `192.168.0.0/16` from managing a bucket named examplebucket.

```
{
    "Version": "1",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "oss:*",
            "Principal": [
                "*"
            ],
            "Resource": [
                "acs:oss:*:174649585760xxxx:examplebucket"
            ],
            "Condition":{
                "NotIpAddress": {
                    "acs:SourceIp": ["192.168.0.0/16"]
                }
            }
        }
    ]
}
```

- Example 3: Allow a RAM user whose UID is `20214760404935xxxx` only to read the `hangzhou/2020` and `hangzhou/2015` directories in a bucket named examplebucket.

```
{
    "Statement": [
        {
            "Action": [
                "oss:GetObject",
                "oss:GetObjectAcl",
                "oss:GetObjectVersion",
                "oss:GetObjectVersionAcl"
            ],
            "Effect": "Allow",
            "Principal": [
                "20214760404935xxxx"
            ],
            "Resource": [
                "acs:oss:*:174649585760xxxx:examplebucket/hangzhou/2020/*",
                "acs:oss:*:174649585760xxxx:examplebucket/hangzhou/2015/*"
            ]
        },
        {
            "Action": [
                "oss:ListObjects",
                "oss:ListObjectVersions"
            ],
            "Condition": {
                "StringLike": {
                    "oss:Prefix": [
                        "hangzhou/2020/*",
                        "hangzhou/2015/*"
                    ]
                }
            },
            "Effect": "Allow",
            "Principal": [
                "20214760404935xxxx"
            ],
            "Resource": [
                "acs:oss:*:174649585760xxxx:examplebucket"
            ]
        }
    ],
    "Version": "1"
}
```

5. Click **Save**.

Access authorized OSS resources

After you configure a bucket policy for a bucket, you can use the following methods to access the resources specified in the bucket policy:

- Object URL (only for authorized anonymous users)

Anonymous users can enter the URL of an object specified in the policy in a browser to access the object. The URL of the object consists of the default domain name of the bucket or a custom domain name mapped to the bucket and the path of the object. Example: `http://mybucket.oss-cn-beijing.aliyuncs.com/file/myphoto.png`. For more information, see OSS domain names.

- OSS console

  Log on to the OSS console. In the left-side navigation pane, click the + icon next to **My OSS Paths**. In the Add Path panel, add the bucket name and the object path specified in the bucket policy. For more information, see Set OSS paths.

- ossutil

  Use the authorized account that is specified in the bucket policy to log on to ossutil to access the resources specified in the policy. For more information, see ossutil.

- ossbrowser

  Use the authorized account that is specified in the bucket policy to log on to ossbrowser. Enter the path of the object specified in the policy in the **Preset OSS Path** field. For more information, see ossbrowser.

- OSS SDK

  You can use OSS SDKs for the following programming languages to access the resources that are specified in the policy: Java, PHP, Node.js, Python, Browser.js, .NET, Android, Go, iOS, C++, and C.

3.7. Download objects

After you upload objects to a bucket, you can download the objects to the default download path of your browser or a specified local path. This topic describes how to download objects from an unversioned bucket in the Object Storage Service (OSS) console.

Prerequisites

When you download Archive or Cold Archive objects, make sure that the objects are restored. For more information, see Restore objects.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the page that appears, click the name of the bucket from which you want to download objects.
3. In the left-side navigation pane, click **Files**, and then download one or more objects.
   - Download a single object

     Method 1: Choose **More > Download** in the Actions column that corresponds to the object that you want to download.

     Method 2: Click the name of the object that you want to download, or click **View Details** in the Actions column that corresponds to the object you want to download. In the **View Details** panel, click **Download**.
   - Download multiple objects

     Select the objects that you want to download. Then, choose **Batch Operation > Download**. You can batch download up to 100 objects in the OSS console.

For more information about how to download objects from versioned buckets, see Configure versioning.

3.8. Configure object metadata

Object Storage Service (OSS) uses object metadata to describe object attributes. Object metadata includes standard HTTP headers and user metadata. You can configure HTTP headers to customize HTTP request policies, such as cache policies and policies for forced object download. You can also configure user metadata to identify the purposes or attributes of objects.

Context

You can configure object metadata for up to 100 objects at a time in the OSS console. To configure object metadata for more than 100 objects at a time, use ossutil.

Procedure

1.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that contains the object for which you want to configure object metadata.
3. In the left-side navigation pane, click **Files**.
4. Use one of the following methods to open the **Set HTTP Header** panel:
    - Configure HTTP headers for one or more objects

      Select one or more objects. Choose **Batch Operation > Set HTTP Header**.
    - Configure HTTP headers for a single object

      Find the object for which you want to configure HTTP headers and choose **More > Set HTTP Header** in the Actions column.
5. In the **Set HTTP Header** panel, configure the parameters. The following table describes the parameters.

Parameter	Description
Content-Type	The type of the object. The browser determines the default method used to open an object based on the object type. For example, the Content-Type value of a GIF image is *image/gif*. For more information about the Content-Type configurations for different object types, see How do I configure the Content-Type of objects?.

Parameter	Description
Content-Encoding	The encoding method of the object. You must set this parameter based on the encoding type of the object. Otherwise, the browser that serves as the client may fail to parse the encoding type of the object, or the object may fail to be downloaded. If the object is not encoded, leave this parameter empty. Default value: identity. Valid values: ○ *identity*: OSS does not compress or encode the object. ○ *gzip*: OSS uses the LZ77 compression algorithm created by Lempel and Ziv in 1977 and 32-bit cyclic redundancy check (CRC) to encode the object. ○ *compress*: OSS uses the Lempel-Ziv-Welch (LZW) compression algorithm to encode the object. ○ *deflate*: OSS uses the zlib library and the deflate algorithm to encode the object. ○ *br*: OSS uses the Brotli algorithm to encode the object. For more information about Content-Encoding, see RFC 2616. **Notice** If you want the static web page objects, such as HTML, JavaScript, XML, and JSON objects to be compressed into GZIP objects when you access these objects, you must leave this parameter empty and add the `Accept-Encoding: gzip` header to your request. For more information, see How do I use GZIP for compression in OSS?
Content-Language	The language of the object content. For example, if the content of an object is written in simplified Chinese, you can set this parameter to *zh-CN*.

Parameter	Description
Content-Disposition	The method used to access the object. Valid values: ○ *inline*: The object is directly opened in the browser. To ensure that an image object or a web page object is previewed but not downloaded when the object is accessed, you must set Content-Disposition to inline and use the custom domain name mapped to the bucket to access the object. For more information about how to map custom domain names, see Map custom domain names. ○ *attachment*: The object is downloaded to the local computer. For example, if this header is set to `attachment; filename="example.jpg"`, the object is downloaded to the local computer. After the object is downloaded, the local file is named `example.jpg`. For more information about Content-Disposition, see RFC 2616.
Cache-Control	The cache configurations for the object. Valid values: ○ *no-cache*: The object can be cached on the client or on the browser of the proxy server. However, each time you access the object, OSS checks whether the cached object is available. If the cache is available, you can directly access the cache. Otherwise, the access request is sent to OSS. ○ *no-store*: All content of the object is not cached. ○ *public*: All content of the object is cached. ○ *private*: All content of the object is cached only on the client. For more information about Cache-Control, see RFC 2616.
Expires	The expiration time of the cache in Greenwich Mean Time (GMT). Example: `2022-10-12T00:00:00.000Z`. If `max-age=<seconds>` is set for Cache-Control, `max-age=<seconds>` takes precedence over Expires.
User Metadata	The user-defined metadata of the object. You can add multiple user metadata headers for an object. However, the total size of user metadata cannot exceed 8 KB. When you add user metadata, user metadata headers must be prefixed with `x-oss-meta-` and assigned values. Example: *x-oss-meta-last-modified:20200909u*.

6. Click **OK**.

3.9. Restore objects

Context

For more information about how to restore objects, see Restore objects.

Procedure

1.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to manage.
3. In the left-side navigation pane, choose **Files > Files**.
4. Find the object that you want to restore and choose ⋮ > **Restore** in the Actions column.

    - Archive objects
        - The restoration takes about one minute. Then, the object is in the restored state.
        - By default, the object remains in the restored state for one day. You can use ossutil or SDKs to extend this period to up to seven days. When the period expires, the object returns to the frozen state.
    - Cold Archive objects

      If you want to restore a Cold Archive object, you must specify the period in which the object can remain in the restored state in **Replica Validity Period**. Unit: days. You can also specify the priority of restoration in **Restore Mode**.

      > **Note** The maximum replica validity period of Cold Archive objects in the China (Shenzhen) region is seven days. The maximum replica validity period of Cold Archive objects outside the China (Shenzhen) region is 365 days.

    The time required to restore an object varies with the size of the object.

3.10. Use OSS Select

The content of an object can be selected and obtained by using simple SQL statements. The amount of data transmitted from Object Storage Service (OSS) can be reduced to improve the data retrieval efficiency by using OSS Select.

Limits

- OSS Select supports UTF-8 encoded CSV files and JSON objects. Supported CSV files and CSV-like files such as TSV files must conform to RFC 4180. You can customize row and column delimiters and quote characters in CSV files.
- You can select and obtain a maximum of 40 MB of data from an object that does not exceed 128 MB in the OSS console. To process an object larger than 128 MB or to retrieve more records, call SelectObject.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket from which you want to select object content.

3. In the left-side navigation pane, choose **Files** > **Files**.
4. On the Files page, choose **More** > **Select Content** in the Actions column corresponding to the object from which you want to select content.
5. In the **Select Content** panel, specify the following parameters:
   - File Type: Select the content type of the object. Valid values: *CSV* and *JSON*.
   - Delimiter: Specify this parameter for CSV files. Valid values: comma (,) and Custom.
   - Title Line: Specify this parameter for CSV files. You can configure this option to specify whether the first row of the object contains a column heading.
   - JSON Display Mode: Select the display mode for JSON objects.
   - Compression Format: Specify whether to compress the current object. Currently, only GZIP-based compression is supported.
6. Click **Preview** to preview the object.

   > **Notice** When you preview Standard objects, you are charged for data scanning by OSS Select. When you preview IA, Archive, and Cold Archive objects, you are charged for data scanning by OSS Select and data retrieval.

7. Click **Next Step**. Enter and execute an SQL statement.
8. View the query result. Click **Download** to download the selected content to the local device.

   For example, a CSV file named *People* contains the following columns: Name, Company, and Age.

   - To query people who are above 50 years old and whose names start with Lora, execute the following SQL statement. In the statement, _1, _2, and _3 specify column indexes. _1 specifies the index of the first column. _2 specifies the index of the second column. _3 specifies the index of the third column.

     ```
     select * from ossobject where _1 like 'Lora*' and _3 > 50
     ```

   - To query the number of rows in the object, maximum age, and minimum age, execute the following SQL statement:

     ```
     select count(*), max(cast(_3 as int)), min(cast(_3 as int)) from oss_object
     ```

3.11. Configure object tagging

You can configure object tagging to classify objects. Object tagging uses key-value pairs to identify objects. You can manage multiple objects that have the same tag. For example, you can configure lifecycle rules for objects that have the same tag.

Context

Object tagging uses key-value pairs to identify objects. The following rules apply to object tagging:

- A maximum of 10 tags can be set for each object. Tags associated with an object must have unique tag keys.
- A tag key can be a maximum of 128 Bytes in length. Each tag value can be a maximum of 256 Bytes in length.
- Keys and values are case-sensitive.
- The key and value of a tag can contain letters, digits, spaces, and special characters such as

+ - = . _ : /

- Only the bucket owner and authorized users have the read and write permissions on object tags. These permissions are independent of the object ACL.
- During cross-region replication (CRR), object tags are also replicated to the destination bucket.

After you configure object tagging, you can configure lifecycle rules for objects that have a specified tag. For example, you can convert the storage class of objects that have a specified tag, or you can delete objects that have a specified tag. For more information, see Configure lifecycle rules. You can authorize Resource Access Management (RAM) users to access objects that have the same tag. For more information, see Object tagging.

Procedure

1. Log on to the OSS console.
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which the objects for which you want to configure object tagging are stored.
3. In the left-side navigation pane, click **Files**.
4. Configure tagging for objects.
   i. Select objects for which you want to configure a tag.
      - Configure tagging for objects in unversioned buckets

        Find the object for which you want to configure tagging, and choose **More > Tagging** in the Actions column.
      - Configure tagging for objects in versioned buckets

        Find the object of a specified version for which you want to configure tagging, and choose **More > Tagging** in the Actions column.

   ii. In the **Tagging** panel, set **Key** and **Value** based on the rules of object tagging.
5. Click **OK**.

3.12. Delete objects

This topic describes how to delete objects that you no longer need from your buckets in Object Storage Service (OSS) by using the OSS console.

Context

You can delete a single object or multiple objects at a time by using the OSS console. You can delete up to 100 objects at a time in the OSS console. Objects can also be selected and deleted in a more flexible manner. For more information about how to delete more than 100 objects at a time, see Delete objects.

Warning

- Objects cannot be restored after they are deleted. Exercise caution when you perform this operation.
-

Procedure

1.

2.
3. In the left-side navigation pane, choose **Files** > **Files**.
4. Select a single object or multiple objects, and choose **Batch Operation** > **Permanently Delete**.
5. In the message that appears, click **OK**.

3.13. Delete directories

You can delete directories from a bucket in the Object Storage Service (OSS) console. When you delete a directory, objects in the directory are also deleted.

Usage notes

- Before you delete a directory, make sure to move objects that you want to retain from the directory to other paths.
- If the directory that you want to delete contains a large number of objects, it can take a long time for OSS to delete the objects. To improve the deletion performance, we recommend that you configure lifecycle rules for the bucket to delete the objects. For more information about how to configure lifecycle rules, see Configure lifecycle rules.
-

Procedure

1.
2.
3. In the left-side navigation pane, choose **Files** > **Files** . Then, delete the directory based on your requirement.

   > **Notice** Do not refresh or close the Task List panel when you delete the directory and the objects in the directory. Otherwise, the tasks are interrupted.

   - Delete a directory from a bucket for which the hierarchical namespace feature is enabled

     Click **Permanently Delete** in the Actions column corresponding to the directory. In the message that appears, click **OK**.

     The directory and the objects in the directory are permanently deleted.

   - Delete a directory from a bucket for which the hierarchical namespace feature is disabled

     If the hierarchical namespace and versioning features are disabled for a bucket, a directory in the bucket and the objects in the directory are deleted in the same way as they are deleted from a bucket for which the hierarchical namespace feature is enabled.

     If the hierarchical namespace feature is disabled and versioning is enabled for the bucket, you can delete the directory and the objects in the directory by using the following methods:

- Store the deleted directory as a previous version

    a. In the upper-right corner of the object list, set **Display Previous Versions** to **Hide**.

    b. Click **Delete** in the Actions column corresponding to the directory. In the message that appears, click **OK**.

    The deleted directory and the objects in the directory are stored as previous versions that can be recovered. For more information about how to recover the previous versions of objects in a directory, see Restore a previous version.

- Permanently delete the directory

    a. In the upper-right corner of the object list, set **Display Previous Versions** to **Show**.

    b. Click **Permanently Delete** in the Actions column corresponding to the directory. In the message that appears, click **OK**.

    The directory and the objects in the directory are permanently deleted.

4. In the **Task List** panel, you can view the deletion progress.

   When OSS performs deletion tasks, you can perform the following operations:

   - Click **Removed** to remove the completed tasks from Task List.
   - Click **Pause All** to pause all running tasks. When tasks are paused, you can perform the following operations:
       - Click **Start** on the right side of a task to resume the task.
       - Click **Remove** on the right side of a task to remove the task. If you remove a paused task, objects that are not deleted in the task are retained.
   - Click **Start All** to resume all paused tasks.

4.Configure OSS DDoS protection

Object Storage Service (OSS) DDoS protection is a proxy-based mitigation service that integrates OSS with Anti-DDoS Pro and Anti-DDoS Premium. When a bucket with OSS DDoS protection enabled suffers DDoS attacks, OSS DDoS protection diverts malicious traffic to an Anti-DDoS instance for scrubbing and then redirects normal traffic to the bucket. This way, your business can continue to function normally after DDoS attacks.

Context

When you use OSS DDoS protection, take note of the following items:

- You can configure OSS DDoS protection only in the China (Hangzhou), China (Shanghai), China (Qingdao), China (Beijing), China (Shenzhen), and China (Hong Kong) regions.
- An Anti-DDoS instance must be retained for at least seven days after the instance is created. If the instance is deleted within seven days, you are charged basic resource fees for the instance for a period of seven days.
- You can create only one Anti-DDoS instance in each region. Each instance can be attached to up to 10 buckets within the same region.

For more information, see OSS DDoS protection.

Procedure

1. Step 1: Create an Anti-DDoS instance.
   i. Log on to the OSS console.
   ii. In the left-side navigation pane, click **Anti-DDoS Pro**.
   iii. On the page that appears, click **Create Anti-DDoS Instance**.
   iv. In the **Create Anti-DDoS Instance** panel, select the required region.
   v. Click **OK**.
2. Step 2: Attach a bucket to the Anti-DDoS instance.
   i. Click **View and Attach Buckets** to the right of the instance to which you want to attach a bucket.
   ii. In the **View and Attach Buckets** panel, click **Attach Buckets**.
   iii. In the **Attach Buckets** panel, select a bucket you want to attach from the **Bucket** drop-down list.

      Buckets that are attached to an Anti-DDoS instance are not displayed in the **Bucket** drop-down list.
   iv. Click **OK**.

      After the bucket is attached to the instance, the bucket is in the **Initializing** status. When the status becomes **Defending**, the Anti-DDoS instance has started to protect the bucket.
3. Step 3: If you want to access the bucket when it is under attack by using the custom domain name that is mapped to the bucket, add the custom domain name in the OSS console.

Notice OSS does not protect the custom domain names mapped to the bucket by default. Therefore, when the bucket is under attack, you cannot access the bucket by using the custom domain names. If you want to access the bucket when it is under attack by using the custom domain names mapped to the bucket, add the custom domain names in the OSS console. You can add up to five custom domain names for each bucket.

- If no custom domain names are mapped to the bucket, you need to map a custom domain name before you add the custom domain name. For more information, see Map custom domain names.
- If a custom domain name is mapped to the bucket, add the custom domain name by performing the following steps:
  a. On the right side of the bucket attached to the instance, choose **More > Modify Custom Domain Name**.
  b. Select the custom domain name that you want to add.
  c. Click **OK**.

     Then, you can access the bucket by using the custom domain name when the bucket suffers attacks.