Object Storage Service

Console User Guide

Console User Guide

The Alibaba Cloud OSS console provides an intuitive operation interface for you to perform most OSS tasks. Before you log on to the OSS console, ensure that you have registered an Alibaba Cloud account. If you do not have an Alibaba Cloud account, the system will prompt you to register an account when you activate OSS.

Operation procedure

Log on to the Alibaba Cloud official website.

On the OSS product detail page, click **Buy now**.

After OSS is activated, click **Console** to access the OSS console.

You can also click **Console** in the upper-right menu bar on the homepage to open Alibaba Cloud console, and click **Object Storage Service** in the left-side navigation pane to access the OSS console.

Manage buckets

All files of Alibaba Cloud OSS are stored in buckets. A bucket is a unit for managing the stored files. All objects must belong to a bucket. You can set the attributes of a bucket for region and file access control and file lifecycle management. These attributes apply to all files in the bucket. Therefore, you can create different buckets to implement different management functions flexibly.

The storage space in a bucket is non-hierarchical, it lacks the features of file systems, such as directories. Therefore, all files are directly affiliated with their corresponding buckets. However, you can group, classify, and manage relevant files by folders.

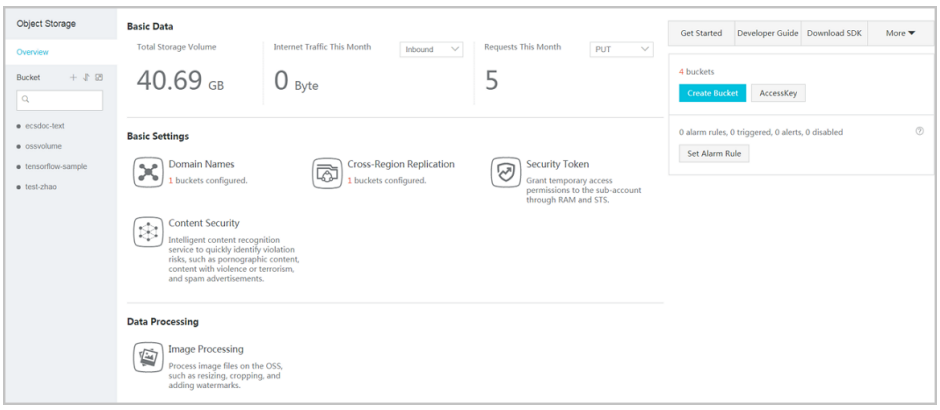

Before uploading any file to the OSS, you must create a bucket to store files. A bucket needs to be configured with various attributes, including its geographic region, access permission, and other metadata.

Procedure

Go to the OSS console.

Click **Create Bucket** or click button **+** in the bucket list to open the **Create Bucket** dialog box.

In the **Bucket Name** text box, enter the bucket name. Note the following when name the bucket. For more information about bucket naming, see Basic OSS concepts.

- The bucket name must comply with the naming rules.
- The bucket name must be unique among all existing bucket names in Alibaba Cloud OSS.
- The bucket name cannot be changed after being created.

In the **Region** drop-down box, select the data center of the bucket.

> **Note**: The region cannot be changed after being subscribed. To access the OSS through the ECS intranet, you can select the same region with your ECS. For more information about regions, see Basic OSS concepts.

In the **Storage Class** drop-down box, select the storage class for the bucket. For the description about storage class, see Introduction to storage classes.

In the **ACL** drop-down box, select an access permission option for the bucket. After creating a bucket, you can set bucket attributes to modify the access permission of the bucket. For

more information about access permission, see Basic OSS concepts.

Private: Only the owner of the bucket can perform read/write operations on the files in the bucket. Other people cannot access the files.

Public Read: Only the owner of the bucket can perform write operations on the files in the bucket, while anyone (including anonymous users) can perform read operations on the files.

Public Read/Write: Anyone (including anonymous users) can perform read and write operations on the files in the bucket. The fees incurred by these operations will be borne by the owner of the bucket. Select this option with caution.

Click **OK** to create the bucket.

If you no longer need a bucket, delete it to avoid paying further fees.

Prerequisite

Before deleting a bucket, make sure that all files in it are cleared, including file fragments caused by incomplete multipart upload. Otherwise, the bucket cannot be deleted.

Note:

- To delete all files in a bucket, see Lifecycle management.
- To delete file fragments, see Manage fragments.

Procedure

Go to the OSS console.

In the bucket name list, choose the bucket that you want to delete.

Click **Delete Bucket** in the upper right corner, then click **OK**.

Note: Buckets cannot be restored after deletion.

The OSS provides an Access Control List (ACL) for permission control. You can configure an ACL when creating a bucket and modify the ACL after creating the bucket. If no ACL is configured, the default value is **Private**.

The OSS ACL provides bucket-level access control. Currently, three access permissions are available for a bucket:

- **Private**: Only the owner of the bucket can perform read/write operations on the files in the bucket. Other people cannot access the files.
- **Public Read**: Only the owner of the bucket can perform write operations on the files in the bucket, while anyone (including anonymous users) can perform read operations on the files.
- **Public Read/Write**: Anyone (including anonymous users) can perform read and write operations on the files in the bucket. The fees incurred by these operations will be borne by the owner of the bucket. Use this permission with caution.

Procedure

Go to the OSS console.

On the left-side navigation pane, select a bucket from the bucket list to open the bucket overview page.

Click the **Basic Settings** tab and find the **ACL** area.

Click **Edit**, and then select an ACL option for the bucket.

Click **Save**.

You can set your bucket to host a static website and access this static website through the bucket domain name.

- If the default webpage is blank, static website hosting is disabled.
- If static website hosting is enabled, we recommend that you use CNAME to bind your domain name.
- Directly accessing the static website root domain or any URL ending with “/ “ under this domain will return the default homepage.

For more information, see Static Website Hosting.

Procedure

Log on to the OSS console.

Select a bucket to open the bucket overview page.

Click **Basic Settings**.

In the **Static Page** area, click **Edit**, then enter the following information:

- **Default Homepage**: The index page (equivalent to the website' s index.html). Only HTML files that have been stored in the bucket can be used.
- **Default 404 Page**: The default 404 page returned when an incorrect path is accessed. Only HTML and image files that have been stored in the bucket can be used. If this field is left empty, the default 404 page is disabled.

Click **Save**.

You can enable or disable logging for a bucket through the console. You can store logs in the same logging-enabled bucket or a new bucket. For more information about the bucket logging format, see Set access logging.

Procedure

Log on to the OSS console.

On the left-side navigation pane, select a bucket from the bucket list to open the bucket overview page.

Click the **Basic Settings** tab and find the **Logs** area.

Click **Edit**, and then edit the logging settings.

- If you do not want to store logs on OSS, close the **Enable Log Storage** switch.

  If you want to store logs on OSS, do as follows:

  a. Open the **Enable Log Storage** switch.

  In the **Log Storage Location** drop-down box, select the name of a bucket to store the logs.

  Note: Only buckets of the same user and region can be selected.

In the **Log Prefix** text box, use the default log prefix **oss-access/**, or click **Custom** to enter another prefix, that is, *<TargetPrefix>* in the following logging naming conventions.

Click **Save**.

Logging naming conventions

The following is the naming conventions for the access log record: *<TargetPrefix><SourceBucket>* YYYY-MM-DD-HH-MM-SS-*<UniqueString>*

- *<TargetPrefix>*: indicates the log prefix specified by the user.
- *<SourceBucket>*: indicates the name of the source bucket.
- YYYY-MM-DD-HH-MM-SS: indicates the time when the log is created. YYYY indicates the year, MM indicates the month, DD indicates the day, HH indicates the hour, MM indicates the minute, and SS indicates the second.
- *<UniqueString>*: indicates the string generated by the OSS.

An example object name used to store OSS access logs is as follows:
MyLog-OSS-example2015-09-10-04-00-00-0000
In the preceding example, **MyLog** is the log prefix specified by the user, **oss-example** is the name of the source bucket, **2015-09-10-04-00-00** is the log creation time, and **0000** is the string generated by the OSS.

The OSS is a Pay-As-You-Go service. To reduce extra fees caused in case your data on the OSS is stolen by others, the OSS supports anti-leech based on the referer field in the HTTP header. You can configure a referer whitelist for a bucket and configure whether to allow the access requests that have an empty referer field.

Procedure

Log on to the **OSS console**.

Select a bucket to open the bucket overview page.

Click **Basic Settings**.

In the **Static Page** area, click **Edit**, then enter the following information:

- **Referer**: Add one or more URLs in whitelist. Separate URLs with carriage returns.
- **Allow Empty Referer**: Whether allow empty referer.

Click **Save**.

Example

For a bucket named test-1-001, set its referer whitelist to http://www.aliyun.com. Then, only requests with a referer of http://www.aliyun.com can access the objects in the bucket.

After uploading an object to a bucket, you can obtain an object address including two parts: an OSS domain name address (*<BucketName>.<Endpoint>*) and an object file name. To avoid possible cross-origin or security problems in your business, we recommend that you access OSS using a user-defined domain name. After the domain name is successfully bound, you also need to add a CNAME record pointing to the Internet domain name of the bucket to ensure proper domain name-based access to the OSS.

Note:

- You must apply for an ICP license for your bound domain name. Otherwise, the domain name is not accessible.
- Each bucket can be bound with a maximum of 20 domain names.

After a user-defined domain name is successfully bound, access addresses of the files stored in your OSS uses the user-defined domain name. For example, if your bucket test-1-001 is located at the Hangzhou node, the object file name is test001.jpg, and the bound user-defined domain name is hello-world.com, then the access address of this object is as follows:

- Before binding: test-1-001.oss-cn-hangzhou.aliyuncs.com/test001.jpg
- After successful binding: hello-world.com/test001.jpg

Bind a domain name

Go to the **OSS console**.

On the left-side navigation pane, select a bucket from the bucket list to open the bucket overview page.

Click the **Domain Names** tab.

Click **Bind User Domain** to open the **Bind User Domain** dialog box.

Bind your domain.

i. In the **User Domain** textbox, enter your domain name.
ii. If you need CDN acceleration, open the **Alibaba Cloud CDN** switch. For details, see CDN-based OSS acceleration.
iii. If you want to add a CNAME record automatically, open the **Add CNAME Record Automatically** switch.

> **Note:** If the domain name has completed cloud resolution under another Alibaba Cloud account, then a CNAME record cannot be automatically added for this domain name under your account. In this case, you must add a CNAME record manually. For details, refer to the **Procedure for domain name resolution** section.

Click **Submit**.

> **Note:** If the domain name you want to bind has been maliciously bound by another user, the system message **Domain name conflict** is displayed. You can verify the ownership of the domain name by adding a TXT record. In this way, the domain name can be forcibly bound to the correct bucket and its binding to the previous bucket is released. For detailed procedure, refer to the **Procedure for verifying domain name ownership** section.

Procedure for verifying domain name ownership

Click **Obtain TXT**. The system generates a TXT record based on your information.

Log on to your DNS provider and add the corresponding TXT record.

In the OSS console, click **I have added the TXT verification file. Continue submission**. If the system detects that the TXT record value for this domain name is as expected, the domain name ownership passes verification.

Procedure for domain name resolution

1. Go to the Alibaba Cloud console.
2. From the left-side navigation pane, click **Alibaba Cloud DNS** to enter the domain name resolution list page.
3. Click the **Configure** link corresponding to the target domain name.
4. Click **Add Record**.
5. In the **Add Record** dialog box, select **CNAME** from the **Type** drop-down box, , and enter the

Internet domain name of the bucket in the **Value** text box.

6. Click **Confirm**.

OSS provides Cross-Origin Resource Sharing (CORS) in the HTML5 protocol to help users achieve cross-origin access. When the OSS receives a cross-origin request (or OPTIONS request), it reads the bucket' s CORS rules and then checks the relevant permissions. The OSS checks each rule sequentially, uses the first rule that matches to approve the request, and returns the corresponding header. If none of the rules match, the OSS does not attach any CORS header.

Procedure

Log on to the **OSS console**.

In the left-side navigation pane, select the target bucket to open the bucket overview page.

Click **Basic Settings**.

In the **Cross-Origin Resource Sharing (CORS)** area, click **Edit**.

In the cross-origin access page, click **Create Rule**.

In the **Cross-Origin Rules** dialog box, configure the following items:

Source: Indicates the origins allowed for cross-origin requests. Multiple matching rules are allowed, which are separated by a carriage return. Only one wildcard () are allowed for each matching rule.

Allowed Methods: Indicates the allowed cross-origin request methods.

Allowed Headers: Indicates the allowed cross-origin request headers. Multiple matching rules are allowed, which are separated by a carriage return. Only one wildcard () are allowed for each matching rule.

Exposed Headers: Indicates the response headers that users are allowed to access from an application (e.g., a Javascript XMLHttpRequest object).

Cache Time: Indicates the cache time for the returned results of browser prefetch (OPTIONS) requests to a specific resource.

Note: A maximum of 10 rules can be configured for each bucket.

Click **OK** to save this rule.

You can define and manage the lifecycle of all the objects or the objects with the same prefix in a bucket. Lifecycle management is used for batch file deletion and automatic fragment deletion.

Note:

- The system will make sure that data is cleared for objects that match a lifecycle rule, within two days from the effective date.
- Configure rules with caution because data deleted based on a lifecycle rule cannot be recovered.

Procedure

Log on to the **OSS console**.

Select a bucket to open the bucket overview page.

Click **Basic Settings**.

In the **Lifecycle** area, click **Edit**.

In the lifecycle page, click **Create Rule**.

In the **Create Lifecycle Rule** dialog box, configure the following:

Status: Specifies whether to enable or disable the rule.

Policy

- **Match by prefix**: If you select this option, you must specify **Prefix**. Then the rule only applies to the objects with the specified prefix, such as **img/**.
- **Apply to bucket**: If you select this option, then the rule applies to all objects in the bucket.

Delete File

Expiration Period: The number of days for retaining an object. If the number of days from the last modification time of the object exceeds the specified number of days, the object will be deleted. For example, if **Expiration Period** is set to **30**, the objects whose last modification date is 2016-1-1 will be scanned and deleted by the backend program on 2016-1-31.

Expiration Date: The objects whose last modification time is earlier than the specified date will be deleted. For example, if **Expiration Date** is set to **2012-12-21**, the objects whose last modification time is earlier than 2012-12-21 will be scanned and deleted by the backend program.

Not Enabled : If this option is selected, files will not be automatically deleted.

Delete Fragments

Expiration Period: The number of days for storing a multipart upload task. If the number of days from the initialization date of a multipart upload task exceeds the specified number of days, the task will be deleted. For example, if **Expiration Period** is set to **30**, the multipart upload tasks whose initialization date is 2016-1-1 will be scanned and deleted by the backend program on 2016-1-31.

Expiration Date: The multipart upload tasks whose last modification time is earlier than the specified date will be deleted. For example, if the **Expiration Date** is set to **2012-12-21**, the multipart upload tasks whose last modification time is earlier than 2012-12-21 will be scanned and deleted by the backend program.

Not Enabled: If this option is selected, fragments will not be automatically deleted.

Click **OK** to save this rule.

Note: After the rule is successfully saved, you can view the configured lifecycle rule in the policy list and perform corresponding **Edit** or **Delete** operations.

Cross-region replication supports synchronization of buckets with different names. If you have two

buckets belonging to different regions, you can enable cross-region replication on the console to synchronize data from the origin bucket to the target bucket in a real-time manner.

Note: Currently, only the regions in mainland China, and the regions between US West 1 and US East 1 support cross-region replication.

Procedure

Log on to the OSS console.

In the left-side navigation pane, select the target bucket to open the bucket overview page.

Click **Basic Settings**.

In the **Cross-Region Replication** area, click **Enable Synchronization**. A dialog box is displayed.

Select **Destination Region** and **Destination Bucket**.

Note:

- Two buckets involved in data synchronization must belong to different regions. Data synchronization cannot be performed between buckets in the same region.
- Two buckets enabled with cross-region replication cannot have a synchronization relationship with any other buckets.

Select **Synchronization Object**.

- **Synchronize All Files**: Synchronize all the files in a bucket to target bucket.
- **Synchronize Files by Prefix**: Synchronize files with specified prefixes in a bucket to target bucket. Up to 10 prefixes can be added.

Select **Synchronization Policy**.

- **Full Synchronization (Add/Delete/Change)**: Synchronize all the data in a bucket to target bucket, including adding, updating, and deleting data.
- **Write Synchronization (Add/Change)**: Synchronize all the data in a bucket to target bucket, including added, changed, and deleted data.

Select **Synchronize History Data**.

Note: When synchronizing history data, files in target buckets which has same names as the files in original buckets may be covered.

Click **OK** to save the setting.

Note: After the configuration is complete, it takes 3 to 5 minutes for cross-region replication to be enabled. Synchronization-related information will be displayed after bucket synchronization.

You can set mirroring rules to define whether to get source data by mirroring or redirection. Mirroring rules are usually used for hot migration of data and redirection of specific requests. You can configure up to five mirroring rules, which is executed by the system in sequence.

Note: Source retrieval does not support intranet endpoint. For more information on traffic fees, see **Pricing**.

Procedure

Log on to the **OSS console**.

Click to open the target bucket.

Click the **Basic Settings** tab.

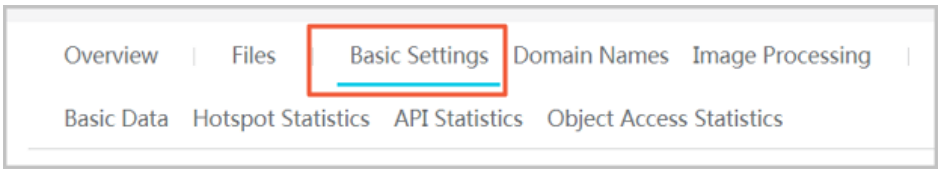

In the **Source Retrieval** area, click **Edit**.

Click **Create Rule**.

Select **Mirroring** or **Redirect** type from the dialog box to configure a rule.

When **Mirroring** is configured, if a requested file is not found on OSS, it will be automatically retrieved from the source site, saved to the OSS, and the content will be returned to the user.

When **Redirect** is configured, the requests that meet the response condition will be returned to the redirected URL through HTTP redirection. A browser or client then obtains the content from the source site.

Click **OK** to save the rule.

Note: After the rule is successfully saved, you can view the configured mirroring rule in the rule list and perform corresponding **Edit** or **Clear** operations.

Manage objects

In OSS, the basic data unit for user operations is an object. The size of a single object is limited to 48.8 TB. An infinite number of objects can exist in a single bucket.

After you create a bucket in a region, the objects uploaded to the bucket are retained in this region, unless you transmit the objects to another region on purpose. Objects stored in an Alibaba Cloud OSS region are physically retained in this region. OSS does not retain copies or move the objects to any other region. However, you can access these objects from anywhere if you have permissions.

You must have the write permission to the bucket before uploading an object to OSS. In the console, the uploaded objects are displayed as files or folders to users. This section describes how to create, manage, and delete files and folders using the console.

After you create a bucket, you can upload all types of files (objects) to the bucket in either of the following ways:

You can upload files smaller than 5 GB by using the OSS console.

You can upload files larger than 5 GB by using SDKs or APIs. For more information, see Multipart upload.

Note: If the name of the file to be uploaded is duplicate with that of the existing file in the bucket, it will overwrite the existing one.

Procedure

Log on to the **OSS console**.

Click to open the target bucket.

Click the **Files** tab.

Click **Upload**.

> **Note:** You can upload a file to a specified folder or the default folder. By clicking Create Directory before clicking **Upload**, you can upload a file to a specified folder. By directly clicking **Upload**, you can upload a file to the OSS default folder.

In the **Directory Address** box, set the path under which the file is uploaded to OSS.

Current Directory: The default path for file uploading. You cannot change the path if selecting this option.

Specify Directory: If you want to upload a file to a certain folder, you must enter the path name. OSS automatically creates the directory and uploads the file to the directory.

> **Note:** For the description of and operations on a folder, see Create a folder.

In the **File ACL** region, select the read/write permissions of the file. The read/write permissions of the bucket where the file belongs are inherited by default.

In the **Upload** region, drag the file to be uploaded to this region, or click **upload them directly** to select the file to be uploaded.

Alibaba Cloud OSS does not has the term **folder**. All elements are stored as objects. To use a folder in the OSS console, you actually create an object with a size of 0 ending with a slash (/) used to sort the same type of files and process them in batches. By default, the OSS console displays objects ending with a slash as folders. These objects can be uploaded and downloaded normally. In the OSS console, you can use OSS folders like using folders in the Windows operating system.

> **Note**: The OSS console displays any object ending with a slash as a folder, whether or not it contains data. The object can be downloaded only using an application programming interface (API) or software development kit (SDK). For more information about how to create and use simulated folders, see API - Get Bucket and **Folder Simulation** in Java SDK- Object.

Procedure

Log on to the OSS console.

Click to open the target bucket.

Select the **Files** tab.

Click **Create Directory**.

Enter a directory name.

Click **OK**.

This section describes how to use the OSS console to search for objects with the same name prefix in a bucket or folder.

When you perform search by name prefix, the search string is case-sensitive and cannot contain the forward slash (/). The search range is limited to the root level of the current bucket or the objects in the current folder (not including subfolders and objects in them). For more information about how to use the forward slash (/) on OSS, see View the object list.

Procedure

Log on to the OSS console.

Click to open the target bucket.

Click **Files**.

Enter the search prefix in the search box, and press **Enter** or click the search icon. The system lists the names of the objects and folders prefixed with **aliyun** in the root directory of the bucket.

To search in a folder, open the folder and enter a search prefix in the search box. The system lists the names of the objects and folders matching the search prefix in the root directory of the folder.

After you upload an object to a bucket, you can get the file address used to share and download the file.

Procedure

Log on to the OSS console.

Click to open the target bucket.

Click the **Files** tab.

Click the name of the target file.

The **Preview** page is displayed.

Copy File URL: used to download the file.

Copy File Path: used to search a file or watermarking an image file.

Click **Copy File URL** and give it to any user who needs to browse or download the file.

If your bucket is set to **Private**, you must set the **Validity** when getting a file URL.

Note: The link validity period for URL signature is calculated based on NTP. You can give this link to any visitor who can then use it to access the file within the validity period. If the bucket has a private permission, the obtained addresses are generated by adding a signature to URL.

You can set an HTTP header for one or multiple files on the OSS console.

You can set an HTTP header for up to 1,000 files using the batch process on the OSS console.

- API: The object header is set through the CopyObject operation.
- SDK: The object header is set through the CopyObject method in Java SDK - Manage objects.

Procedure

Log on to the OSS console.

Click to open the target bucket.

Click **Files**.

Select one or multiple files, and then click **Set HTTP Header**.

You can also click one file name and then click **Set HTTP Header** on the **Preview** page.

Enter the values. For more information about each field, see Definitions of common HTTP headers.

Click **OK**.

If you do not need to store uploaded files any longer, delete them to avoid further fees. You can delete a single file or multiple files on the OSS console.

Note: The deleted file cannot be recovered. Perform this operation with caution.

You can delete up to 1,000 files at a time on the console. If you want to delete only the selected files or perform batch deletion in a larger volume, follow the procedures in API or SDK documents. For more information, see the relevant sections of the Developer Guide.

- API: Delete Object and Delete Multiple Objects
- SDK: **Delete multiple objects** in Java SDK - Manage objects

Procedure

Log on to the OSS console.

Click to open the target bucket.

Click **Files**.

Select one or multiple files, and then click **Delete**.

Click **OK**.

After you delete a folder on the OSS console, all files and sub folders in this folder are automatically

deleted. If you want to retain the files, move them to other places before you delete the folder.

Procedure

Log on to the OSS console.

Click to open the target bucket.

Click **Files**.

Select the target folder, and then click **Delete**.

Note: Deletion may fail if the folder contains too many files.

Click **OK** to delete the folder.

In multipart upload, an object is split into several parts. These parts are uploaded separately. When all parts are uploaded, the OSS API CompleteMultipartUpload is called to combine the parts into the complete object. If multipart upload fails, unuseful parts, called fragments, may be accumulated in OSS buckets. In this case, we recommend that you delete these fragments to save your storage space on OSS. Besides, if you want to delete a bucket, you must delete all the objects and fragments in the bucket first.

Procedure

1. Log on to the OSS console.
2. On the left-side navigation pane, select a bucket from the bucket list to open the bucket overview page.
3. Click the **Files** tab.
4. Click **Fragments**.
5. In the **Fragments** dialog box, delete fragments.
   - If you want to delete all the fragments in the bucket, click **Clear All Fragments**.
   - If you want to delete certain fragments, select or search for the fragments, and then click **Delete Fragments**.
6. In the **Clear Fragments** dialog box, click **OK** to confirm the deletion.