ALIBABA CLOUD

Alibaba Cloud

Apsara Stack Enterprise

Object Storage Service

User Guide

Product Version: V3.16.2

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.What is OSS?

Object Storage Service (OSS) is a secure, cost-effective, and highly reliable cloud storage service provided by Alibaba Cloud.

Compared with user-created server storage, OSS has outstanding advantages in reliability, security, cost-effectiveness, and data processing capabilities. OSS enables you to store and retrieve a variety of unstructured data objects, such as text, images, audios, and videos over networks anytime.

OSS is an object storage service based on key-value pairs. Files uploaded to OSS are stored as objects in buckets. You can obtain the content of an object based on the object key.

In OSS, you can perform the following operations:

- Create a bucket and upload objects to the bucket.
- Obtain an object URL from OSS to share or download the object.
- Modify the attributes or metadata of a bucket or an object. You can also configure the access control list (ACL) of the bucket or the object.
- Perform basic and advanced operations in the OSS console.
- Perform basic and advanced operations by using OSS SDKs or calling RESTful API operations in your application.

2.Usage notes

Before you use OSS, you must understand the following content:

To allow other users to use all or part of OSS features, you must create RAM users and grant permissions to the users by configuring RAM policies.

Before you use OSS, you must also understand the following limits.

Item	Limit
Bucket	• You can create up to 100 buckets. • After a bucket is created, its name and region cannot be modified.
Upload objects	• Objects larger than 5 GB cannot be uploaded by using the following modes: console upload, simple upload, form upload, or append upload. To upload an object that is larger than 5 GB, you must use multipart upload. The size of an object uploaded by using multipart upload cannot exceed 48.8 TB. • If you upload an object that has the same name of an existing object in OSS, the new object will overwrite the existing object.
Delete objects	• Deleted objects cannot be recovered. • You can delete up to 100 objects at a time in the OSS console. To delete more than 100 objects at a time, you must call an API operation or use an SDK.
Lifecycle	You can configure up to 1,000 lifecycle rules for each bucket.

3.Quick start

3.1. Log on to the OSS console

This topic describes how to log on to the Object Storage Service (OSS) console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

   Obtain the username and password that you can use to log on to the console from the operations administrator.

   > **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 10 to 32 characters in length and must contain at least two of the following character types:
   >
   > - Uppercase or lowercase letters
   > - Digits
   > - Special characters, which include ! @ # $ %

3. Click **Log On**.
4. If your account has multi-factor authentication (MFA) enabled, perform corresponding operations in the following scenarios:
   - It is the first time that you log on to the console after MFA is forcibly enabled by the administrator.
     a. On the Bind Virtual MFA Device page, bind an MFA device.
     b. Enter the account and password again as in Step 2 and click **Log On**.
     c. Enter a six-digit MFA verification code and click **Authenticate**.
   - You have enabled MFA and bound an MFA device.

     Enter a six-digit MFA authentication code and click **Authenticate**.

     > **Note** For more information, see the *Bind a virtual MFA device to enable MFA* topic in *Apsara Uni-manager Operations Console User Guide*.

5. In the top navigation pane, choose **Products > Object Storage Service**.

3.2. Create buckets

Objects uploaded to OSS are stored in a bucket. You must create a bucket before you upload objects to OSS.

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. On the **Create OSS Bucket** page, configure parameters.

   The following table describes the parameters that you can configure.

Parameter	Description
Organization	Select an organization from the drop-down list for the bucket.
Resource Set	Select a resource set from the drop-down list for the bucket.
Region	Select a region from the drop-down list for the bucket. **Note** ◦ The region of a bucket cannot be changed after the bucket is created. ◦ If you want to access OSS from your ECS instance through the internal network, select the same region where your ECS instance is deployed.
Cluster	Select a cluster for the bucket.
Bucket Name	Enter the name of the bucket. **Note** ◦ The bucket name must comply with the naming conventions. ◦ The bucket name must be globally unique among all existing buckets in OSS. ◦ The bucket name cannot be changed after the bucket is created.

Parameter	Description
Storage Class	Set the value to **Standard**. Only Standard is supported.
Bucket Capacity	Specify the capacity of the bucket. Valid values: 0 to 2000000. Unit: TB or GB.
Access Control List (ACL)	Set the ACL of the bucket. You can select the following options: ○ **Private**: Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users cannot access objects in the bucket without authorization. ○ **Public Read**: Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users can only read objects in the bucket. ○ **Public Read/Write**: All users, including anonymous users can read and write objects in the bucket. Fees incurred by such operations are paid by the owner of the bucket. Exercise caution when you configure this option. **Note** You can modify the ACL of a bucket after the bucket is created. For more information, see Modify bucket ACLs.
Server-Side Encryption	Configure server-side encryption for the bucket. You can select the following options: ○ **None**: Server-side encryption is not performed. ○ **AES256**: AES256 is used to encrypt each object in the bucket using different data keys. Customer master keys (CMKs) used to encrypt the data keys are rotated regularly. ○ **KMS**: CMKs managed by KMS are sued to encrypt objects in the bucket.
Encryption Algorithm	You can configure this parameter when you select **KMS** for **Server-Side Encryption**.
Key ID	You can configure this parameter when you select **KMS** for **Server-Side encryption**. OSS uses the specified CMK to encrypt objects in the bucket. **Note** To select a CMK ID for server-side encryption, you must create the CMK in the KMS console.

4. Click **Submit**.

3.3. Upload objects

After you create a bucket, you can upload objects to it.

Prerequisites

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

Context

You can upload an object of any format to a bucket. You can use the OSS console to upload an object up to 5 GB in size. To upload an object larger than 5 GB, use OSS SDKs or call an API operation.

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. On the bucket details page that appears, click **Files**.
4. Click **Upload**.
5. In the **Upload** panel, set the parameters described in the following table.

Parameter	Description
Upload To	Set the directory to which you want to upload objects. ○ **Current**: Objects are uploaded to the current directory. ○ **Specified**: Objects are uploaded to the specified directory. You must enter the directory name. If the specified directory does not exist, OSS automatically creates the specified directory and uploads the object to the directory.
File ACL	Set the access control list (ACL) of the object to upload. Default value: **Inherited from Bucket**. ○ **Inherited from Bucket**: The ACL of uploaded objects is the same as that of the bucket. ○ **Private**: Only the owner or authorized users can read and write objects in the bucket. Other users, which includes anonymous users, cannot access the objects in the bucket without authorization. ○ **Public Read**: Only the bucket owner can perform write operations on objects in the bucket. Other users, which includes anonymous users, can perform only read operations on objects in the bucket. ○ **Public Read/Write**: All users, which includes anonymous users, can read and write objects in the bucket.

Parameter	Description
Upload	Drag one or more files to upload to this section, or click **Upload** to select one or more files to upload. **Notice** ◦ When the object to upload has the same name as an existing object in the bucket, the existing object is overwritten. ◦ If you upload a directory, only the files in the directory are uploaded, and the files are stored in the same directory in the bucket. ◦ The name of an uploaded object must comply with the following conventions: ▪ The name must be encoded in UTF-8. ▪ The name is case-sensitive. ▪ The name must be 1 to 1,023 bytes in length. ▪ The name cannot start with a forward slash (/) or backslash (\).

6. In the **Upload Tasks** panel, wait until the upload task is completed.

   During the upload process, you can click **Cancel All** to cancel the task. After the task is completed, click **Removed** to remove the task.

   Notice Do not refresh or close the **Upload Tasks** panel when objects are being uploaded. Otherwise, the upload tasks are interrupted.

3.4. Obtain object URLs

You can obtain the URL of an uploaded object in Object Storage Service (OSS) and share the URL with other users to preview or download the object.

Prerequisites

An object is uploaded to the bucket. For more information, see Upload objects.

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click the **Files** tab.
4. Obtain object URLs.
   - Obtain the URL of a single object.

     Click the name of the object whose URL you want to obtain, or click **View Details** in the Actions column that corresponds to the object. In the **View Details** panel, click **Copy File URL**.
   - Batch export URL lists.

     Select the objects that you want to share. Choose **Batch Operation > Export URL List**.

4.Buckets

4.1. View bucket information

You can view the detailed information about created buckets in the Object Storage Service (OSS) console.

Prerequisites

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

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. On the **Overview** tab, you can view the information about the bucket, which includes Organization and Resource Set, Domain Names, and Basic Settings.

4.2. Delete buckets

You can delete a bucket in the Object Storage Service (OSS) console.

Prerequisites

All objects and parts in the bucket are deleted. For more information, see Delete objects and Manage parts.

> **Warning** Deleted objects, parts, and buckets cannot be recovered. Exercise caution when you delete objects, parts, and buckets.

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. In the upper-right corner, click **Delete Bucket**.
4. In the message that appears, click **OK**.

4.3. Modify bucket ACLs

Object Storage Service (OSS) provides access control list (ACL) to control access to buckets. By default, the ACL of a bucket is private. You can modify the ACL of a bucket after the bucket is created.

Prerequisites

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

Context

You can set the ACL of a bucket to one of the following values:

- Private: Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users, cannot access the objects in the bucket without authorization.
- Public Read: Only the owner or authorized users of this bucket can write objects in the bucket. Other users, including anonymous users, can only read the objects in the bucket.
- Public Read/Write: Any users, including anonymous users, can read and write the objects in the bucket.

Warning If you set the ACL of a bucket to Public Read or Public Read/Write, other users can read the data in the bucket without authentication, which may result in security risks. To ensure the security of your data, we recommend that you set the ACL of your bucket to private.

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click the **Basic Settings** tab. Find the **Access Control List (ACL)** section.
4. Click **Configure**. Modify the bucket ACL.
5. Click **Save**.

4.4. Configure static website hosting

You can configure static website hosting for a bucket in the Object Storage Service (OSS) console so that users can access the website by using the domain name of the bucket.

Prerequisites

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

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click the **Basic Settings** tab. Find the **Static Pages** section.
4. Click **Configure** and then set the parameters described in the following table.

Parameter	Description
Default Homepage	Specify an index page that functions similar to index.html. Only HTML objects can be specified as the index page. If you do not specify this parameter, static website hosting is disabled.
Default 404 Page	Set the default 404 page that is displayed when the requested resource does not exist. Only HTML, JPG, PNG, BMP, or WebP objects in the root directory of the bucket can be set to the default 404 page. If you do not specify this parameter, Default 404 Page is disabled.

5. Click **Save**.

4.5. Configure hotlink protection

You can configure hotlink protection for a bucket in the Object Storage Service (OSS) console to prevent data in your bucket from being accessed by unauthorized domain names.

Prerequisites

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

Context

The hotlink protection feature allows you to configure a Referer whitelist for a bucket. This way, only requests from domain names included in the Referer whitelist can access your data in the bucket. OSS allows you to configure Referer whitelists based on the Referer header field in HTTP or HTTPS requests.

After hotlink protection is configured for a bucket, OSS verifies requests to objects in the bucket only when the requests are initiated by using signed URLs or from anonymous users. Requests that contain the Authorization field in the header are not verified.

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click the **Basic Settings** tab. Find the **Hotlink Protection** section.
4. Click **Configure** and configure the following parameters:
   - Enter domain names or IP addresses in the **Referer Whitelist** field. Separate multiple Referers by using line feed. You can use asterisks (*) and question marks (?) as wildcards. Examples:
     - If you add www.example.com to the Referer whitelist, requests sent from URLs that start with *www.example.com*, such as *www.example.com/123* and *www.example.com.cn* are allowed.
     - If you add *www.example.com/ to the Referer whitelist, requests sent from *http://www.example.com/* and *https://www.example.com/* are allowed.
     - An asterisk (*) can be used as a wildcard to indicate zero or more characters. For example, if you add *.example.com to the Referer whitelist, requests sent from URLs such as *help.example.com* and *www.example.com* are allowed.
     - A question mark (?) can be used as a wildcard to indicate a single character. For example, if you add example?.com to the Referer whitelist, requests sent from URLs such as *examplea.com* and *exampleb.com* are allowed.
     - 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 in which the Referer field is empty.

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

     If you do not allow empty Referers fields, only HTTP or HTTPS requests which include an allowed Referer field can access the objects in the bucket.

> **Note** By default, if you use the bucket endpoint to preview an MP4 object, the browser sends a request that contains the Referer field and a request that does not contain the Referer field at the same time. Therefore, to allow access to the MP4 objects in your bucket, you must not only add the bucket endpoint to the Referer whitelist but also allow empty Referer fields. To preview a non-MP4 object by using the bucket domain name, you need only to allow empty Referer fields.

5. Click **Save**.

4.6. Configure CORS

You can configure cross-origin resource sharing (CORS) in the Object Storage Service (OSS) console to enable cross-origin access.

Prerequisites

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

Context

OSS provides CORS over HTML5 to implement cross-origin access. When OSS receives a cross-origin request (or an OPTIONS request) for a bucket, OSS reads the CORS rules of the bucket and checks the relevant permissions of the request. OSS matches the request with the rules one by one. When OSS finds the first match, OSS returns a corresponding header in the response. If no match is found, OSS does not include any CORS header in the response.

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click the **Basic Settings** tab. In the **Cross-Origin Resource Sharing (CORS)** section, click **Configure**.
4. On the page that appears, click **Create Rule**. Then, in the **Create Rule** panel, configure the parameters in the following table.

Parameter	Required	Description

Parameter	Required	Description
Sources	Yes	Specify the sources from which you want to allow cross-origin requests. Take note of the following items when you configure the sources: ○ 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. ○ Asterisks (*) are supported as wildcards. Each rule can contain up to one asterisk (*). ○ A domain name must include the port number if the domain name does not use the default port. 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. ○ Use an asterisk (*) as a wildcard in the domain name to match second-level domains. Example: https://*.example.com. ○ Enter only an asterisk (*) as the wildcard to match all domain names.
Allowed Methods	Yes	Select the cross-origin request methods that are allowed.
Allowed Headers	No	Specify the response headers for the allowed cross-origin requests. Take note of the following rules when you configure the allowed response headers: ○ This parameter is in the key:value format and case-insensitive. Example: content-type:text/plain. ○ You can configure multiple rules for allowed headers. Separate multiple rules with new lines. ○ Each rule can contain up to one asterisk (*) as the wildcard. We recommend that you set this parameter to an asterisk (*) if you do not have special requirements.
Exposed Headers	No	Specify the response headers for allowed access requests from applications, such as an XMLHttpRequest object in JavaScript. Exposed headers cannot include asterisks (*).

Parameter	Required	Description
Cache Timeout (Seconds)	No	Specify the period of time within which the browser can cache the response for an OPTIONS preflight request to a specific resource. Unit: seconds.

> **Note** You can configure up to 10 CORS rules for each bucket.

5. Click **OK**.

4.7. Configure lifecycle rules

You can configure a lifecycle rule for a bucket to regularly delete expired objects and parts to save storage costs.

Prerequisites

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

Context

Take note of the following items when you configure lifecycle rules for a bucket:

- After a lifecycle rule is configured, it is loaded within 24 hours and takes effect within 24 hours after it is loaded. Check the configurations of a rule before you save the rule.
- Objects that are deleted based on lifecycle rules cannot be recovered. Configure lifecycle rules based on your requirements.
- You can configure up to 100 lifecycle rules for each bucket in the Object Storage Service (OSS) console and up to 1,000 lifecycle rules for each bucket by using ossutil.

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 lifecycle rules.
3. Click the **Basic Settings** tab. Find the **Lifecycle** section. Click **Configure**.
4. Click **Create Rule**. In the **Create Rule** panel that appears, configure the parameters described in the following table.

Parameter	Description
Status	Specify the status of the lifecycle rule. Valid values: **Enabled** and **Disabled**.

Parameter	Description
Applied To	Select objects to which the rule applies. You can select **Files with Specified Prefix** or **Whole Bucket**. Files with Specified Prefix indicates that this rule applies to objects whose names contain a specified prefix. Whole Bucket indicates that this rule applies to all objects in the bucket. **Note** If you select **Files with Specified Prefix**, you can configure different lifecycle rules for objects whose names contain different prefixes. If you select **Whole Bucket**, only one lifecycle rule can be configured for the bucket.
Prefix	If you set **Applied To** to **Files with Specified Prefix**, you must specify the prefix of the objects to which the rule applies. For example, if you want that the rule applies to objects whose names start with img, enter *img*.
File Lifecycle	Configure rules for objects to specify when objects expire. You can set File Lifecycle to **Validity Period (Days)**, **Expiration Date**, or **Disabled**. If you select **Disabled**, the configurations of File Lifecycle do not take effect.
Delete	Specify when objects expire based on **Validity Period (Days)** or **Expiration Date** that you set for **File Lifecycle**. After objects expire, the objects are deleted. ○ **Validity Period (Days)**: Specify the number of days to retain objects after they are last modified. Select the check box next to Delete and specify a number such as N. The objects expire N days after they are last modified. Then, the objects are deleted the next day after they expire. For example, if you specify the number as 30, objects that are last modified on January 1,2019 are deleted on February 1, 2019. ○ **Expiration Date**: Specify the expiration date. Select the check box next to Delete and set an expiration date. The objects that are last modified before this date expire and are deleted. For example, if you set Expiration Date to January 1, 2019, objects that are last modified before January 1, 2019 are deleted.
Part Lifecycle	Specify the operations to perform on expired parts. You can set Part Lifecycle to **Validity Period (Days)**, **Expiration Date**, or **Disabled**. If you select **Disabled**, the configurations of Part Lifecycle do not take effect. **Notice** You must configure at least one of **File Lifecycle** and **Part Lifecycle**.
Delete Parts	Specify when parts that match the rule expire based on **Validity Period (Days)** or **Expiration Data** that you set for **Part Lifecycle**. Expired parts are deleted. You can configure this parameter in the same way as you configure the Delete parameter in Clear Policy.

5. Click **OK**.

4.8. Configure storage quota

If the capacity of a bucket reaches the specified storage quota, write operations such as PutObject, MultipartUpload, CopyObject, PostObject, and AppendObject cannot be performed on the bucket. This topic describes how to configure the storage quota of a bucket in Object Storage Service (OSS).

Prerequisites

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

Context

Take note of the following items when you configure the storage quota of a bucket:

- Before you configure the storage quota of a bucket, make sure that the quota does not limit your business because write operations cannot be performed if the bucket capacity reaches the specified quota.
- In general, it takes about an hour for OSS to determine whether the bucket capacity exceeds the storage quota. In some cases, it can take longer than one hour.

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 storage quota.
3. Click the **Basic Settings** tab and find the **Storage Quota** section.
4. Click **Configure**.
5. On the **Modify OSS Bucket** page, modify the storage quota of the bucket.
   - Units: TB or GB.
   - Valid values: -1 to 2000000

     The default value is -1, which indicates that the bucket capacity is not limited.
6. Click **Submit**.

   After you submitted, you can click **Back to Console** in the pop-up dialog box to go back to the **Overview** page.

4.9. Configure cluster-disaster recovery

In cluster-disaster recovery mode, buckets with the same name are replicated. Cluster-based disaster recovery is automatically enabled based on configurations made when the cluster is created. In other words, after a primary bucket is created, a secondary bucket with the same name is automatically created. Information stored in the primary bucket is automatically synchronized to the secondary bucket. By default, Cluster-disaster Recovery is turned on for buckets that are created by using the Object Storage Service (OSS) console.

Prerequisites

A bucket is created. For more information about how to create a bucket, see Create buckets.

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 go to the bucket details page.
3. On the bucket details page, click the **Basic Settings** tab. Find the **Cluster-disaster Recovery** section.
4. Click **Configure**. Turn on or turn off **Cluster-disaster Recovery**.
5. Click **Save**.

4.10. Bucket tagging

Object Storage Service (OSS) allows you to configure bucket tagging to classify and manage buckets. For example, you can use this feature to list buckets that have specific tags and configure access control lists (ACLs) for buckets that have specific tags.

Context

The bucket tagging feature uses a key-value pair to identify a bucket. You can add tags to buckets that are used for different purposes and manage the buckets by tags.

- Only the bucket owner or authorized Resource Access Management (RAM) users can configure tagging for the bucket. Otherwise, 403 Forbidden is returned with the AccessDenied error code.
- You can configure up to 20 tags for a bucket.
- The tag key is required. The tag key can be up to 64 Bytes in length and cannot start with `http://`, `https://`, or `Aliyun`.
- The tag value is optional. The tag value can be up to 128 bytes in length. You can leave the parameter empty.
- The key and value of a tag must be encoded in UTF-8.

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 tagging.
3. Choose **Basic Settings > Bucket Tagging**.
4. Click **Configure**.
5. Add tags to the bucket based on the naming conventions. You can click the **+** icon to add multiple tags to a bucket.
6. Click **Save**.

4.11. Configure server-side encryption

OSS supports server-side encryption. When you upload an object to a bucket for which server-side encryption is enabled, OSS encrypts the object and stores the encrypted object. When you download the encrypted object from OSS, OSS automatically decrypts the object and returns the decrypted object to you. A header is added in the response to indicate that the object is encrypted on the OSS server.

Context

OSS supports the following encryption methods:

Context

- Server-side encryption by using KMS (SSE-KMS)

  OSS uses the default customer master key (CMK) managed by KMS or a specified CMK to encrypt objects. The CMK is managed by KMS to ensure confidentiality, integrity, and availability at minimal costs.

- Server-side encryption by using OSS-managed keys (SSE-OSS)

  OSS uses data keys to encrypt objects and manages the data keys. In addition, OSS uses master keys that are regularly rotated to encrypt data keys.

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 on the Basic Settings tab

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. On the **Create OSS Bucket** page, set parameters.

   You can set the following parameters to configure server-side encryption for the bucket.

   - **Server-Side Encryption**: Specify the encryption methods.
     - **None**: Server-side encryption is not performed.
     - **AES256**: AES256 is used to encrypt each object in the bucket by using different data keys. The CMKs used to encrypt the data keys are rotated regularly.
     - **KMS**: CMKs managed by KMS are used to encrypt objects in the bucket.
   - **Encryption Algorithm**: This parameter can be configured when you select **KMS** for **Server-Side Encryption**.
   - **Key ID**: This parameter can be configured when you select **KMS** for **Server-Side Encryption**. OSS uses the specified CMK to encrypt objects in the bucket.

     > **Note**
     >
     > To select a CMK ID for server-side encryption, you must create the CMK in the KMS console.

4. Click **Submit**.

Method 2: Enable server-side encryption on the Basic Settings tab

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 server-side encryption.
3. Click the **Basic Settings** tab. Find the **Server-side Encryption** section.
4. Click **Configure** and set the following parameters:

- **Encryption Method**: Specify the encryption method.
  - **None**: Server-side encryption is not performed.
  - **OSS-Managed**: Keys managed by OSS are used to encrypt your data.
  - **KMS**: CMKs managed by KMS are used to encrypt objects in the bucket.
- **CMK**: This parameter can be configured when you select **KMS** for **Encryption Method**. OSS uses the specified CMK to encrypt objects in the bucket.

> **Note**
>
> To select a CMK ID for server-side encryption, you must create the CMK in the KMS console.

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.

4.12. Bind a bucket to a VPC network

You can bind your bucket to a specified virtual private cloud (VPC) network to allow only requests from IP addresses within the VPC network to access your bucket.

Prerequisites

A VPC network is created. For more information, see the "Create a VPC" chapter of *VPC User 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 bind to the VPC network.
3. Click the **Overview** tab. Click **Bind VPC** in the **VPC Info** section.
4. On the **Bind VPC** page, select the VPC network that you create.

   You can also click **Create VPC** to create a new VPC network.
5. Click **Submit**.

4.13. Configure CRR

Cross-region replication (CRR) allows you to perform automatic and asynchronous (near real-time) replication on objects across buckets that are located in different Object Storage Service (OSS) regions. If you enable CRR, operations such as the creation, overwriting, and deletion of objects can be synchronized from the source bucket to the destination bucket.

Prerequisites

The source bucket and destination bucket are created. For more information, see Create buckets.

Context

CRR meets the requirements of geo-disaster recovery or data replication. Objects in the destination bucket are extra duplicates of objects in the source bucket. They have the same names, content, and metadata, such as the created time, owner, user metadata, and object access control list (ACL).

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 CRR.
3. On the bucket details page, click the **Basic Settings** tab. Find the **Cross-Region Replication** section.
4. Click **Enable**. In the **Cross-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 Region	Select the region in which the destination bucket is located. The source and destination buckets for CRR must be located in different regions. Data cannot be synchronized between buckets that are located within the same region.
Destination Bucket	Select the destination bucket to which data is synchronized. The source bucket and destination bucket specified in a CRR rule are not allowed to synchronize data with other buckets. For example, if you configure a CRR rule to synchronize data from Bucket A to Bucket B, Bucket A and Bucket B are not allowed to synchronize data with other buckets.
Applied To	Select the source data that you want to synchronize. ◦ **All Files in Source Bucket**: All objects within the source bucket are synchronized to the destination bucket. ◦ **Files with Specified Prefix**: Only objects whose names contain one of the specified prefixes are synchronized to the destination bucket. For example, if you have a directory named *management* in the root directory of a bucket and want to synchronize objects in a subdirectory named *abc* in *management*, you can enter the prefix *management/abc*. You can specify up to 10 prefixes.
Operations	Select the synchronization policy. ◦ **Add/Change**: Only the added or updated data is synchronized from the source bucket to the destination bucket. ◦ **Add/Delete/Change**: All changes to data including the creation, modification, and deletion of objects are synchronized from the source bucket to the destination bucket.

Parameter	Description
Replicate Historical Data	Specify whether to synchronize historical data that exists 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 destination bucket, which have the same name as objects from the source bucket, may be overwritten. ○ **No**: Only objects that are uploaded or updated after CRR is enabled are synchronized to the destination bucket.

5. Click **OK**.

Note

- It takes 3 to 5 minutes for a CRR rule to take effect after the rule is configured. Synchronization information is displayed after the source data is synchronized to the destination bucket.
- In CRR, data is asynchronously replicated in near real-time. It may take a few minutes to several hours to replicate data to the destination bucket based on the data amount.

4.14. Configure cross-cloud replication

You can use the cross-cloud replication feature to synchronize Object Storage Service (OSS) data between two clouds. This topic describes how to configure cross-cloud replication.

Step 1: Obtain the parameters of the destination cloud.

Before you configure cross-cloud replication, you must obtain the required parameters of the destination cloud.

1. Log on to the Apsara Uni-manager Operations Console of the destination cloud.

   For more information about how to log on to the Apsara Uni-manager Operations console, see *Log on to the Apsara Uni-manager Operations console* in *Operations and Maintenance Guide*.

2. In the left-side navigation pane, choose **Product Management > Products**.
3. Click **OSS O&M**.
4. In the left-side navigation pane, choose **Service O&M - OSS > Synchronization Management > Cross-Cloud Synchronization**.
5. In the upper-right corner, select the destination cluster from the **Cluster** drop-down list, and then click **View Parameters of the Current Cloud**.

   Record the information displayed in the **Parameters of the Current Cloud** dialog box.

Step 2: Configure cross-cloud synchronization for the source cloud.

After you obtain the required parameters of the destination cloud, you must configure cross-cloud synchronization for the source cloud in the Apsara Uni-manage Operations Console.

1. Log on to the Apsara Uni-manager Operations Console of the source cloud.
2. In the left-side navigation pane, choose **Product Management > Products**.
3. Click **OSS O&M**.
4. In the left-side navigation pane, choose **Service O&M - OSS > Synchronization Management > Cross-Cloud Synchronization**.
5. In the upper-right corner, click **Create**. In the **Create Cross-Cloud Synchronization Task** dialog box, add the obtained parameters of the destination cloud.
6. Click **Submit**.

   Wait a few minutes until the cross-synchronization configurations take effect.

Step 3: Configure cross-cloud replication in the Apsara Uni-manager Management Console of the source cloud.

After the cross-cloud synchronization configurations take effect, you must configure cross-cloud replication in the Apsara Uni-manager Management Console.

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 cross-cloud replication.
3. On the bucket details page, click the **Basic Settings** tab. In the **Cross-Cloud Replication** section, click **Enable**.
4. In the **Cross-Cloud Replication** panel, configure the parameters described in the following table.

Parameter	Description
Source Region	The region in which the source bucket is located.
Source Bucket	The name of the source bucket.
Destination Cloud	Enter the name of the destination cloud obtained in Step 1.
Destination Cloud Address	Enter the value of Location of the destination cloud obtained in Step 1.
Destination Bucket	Enter the name of the destination bucket.
Applied To	Select the objects that you want to synchronize. ○ **All Files in Source Bucket**: All objects within the source bucket are synchronized to the destination bucket. ○ **Files with Specified Prefix**: Only objects whose names contain one of the specified prefixes are synchronized to the destination bucket. Click **Add**. You can add up to 10 prefixes.

Parameter	Description
Operations	Select a synchronization policy. ○ **Add/Change**: Only newly added and changed data is synchronized from the source bucket to the destination bucket. ○ **Add/Delete/Change**: All changes to data including creation, modification, and deletion of objects are synchronized from the source bucket to the destination bucket.
Replicate Historical Data	Specify whether to synchronize historical data that is generated before you enable cross-cloud replication. ○ **Yes**: Historical data is synchronized to the destination bucket. ○ **No**: Only objects that are uploaded or updated after cross-cloud replication is enabled are synchronized to the destination bucket.

5. Click **OK** to save your settings.

4.15. Log management

4.15.1. Configure logging

When you access Object Storage Service (OSS), a large number of access logs are generated. You can use the logging feature to store OSS access logs in a specified bucket.

Prerequisites

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

Procedure

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click the **Basic Settings** tab. Find the **Logging** section.
4. Click **Configure**. Turn on the **Logging** switch. Select **Destination Bucket** and set **Log Prefix**.
   - **Destination Bucket**: Select the bucket in which you want to store access logs from the drop-down list. You must be the owner of the selected bucket, and the selected bucket must be in the same region as the bucket for which logging is enabled.
   - **Log Prefix**: Enter the prefix and directory where the access logs are stored. If you specify *log/<TargetPrefix>* as the prefix, access logs are stored in the *log/* directory.
5. Click **Save**.

4.15.2. Real-time log query

A large number of logs are generated when Object Storage Service (OSS) resources are accessed. 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.

Benefits

- Pushes logs to Log Service within three minutes and allows you to view real-time logs in the OSS console.
- Provides log analysis and common analysis reports so that you can easily query data.
- Allows you to query and analyze raw logs in real time and filter logs by bucket name, object name, API operation, or time.

Prerequisites

Log Service is activated and is authorized to access OSS.

Enable real-time log query

> **Notice**
>
> - After you enable real-time log query for a bucket for the first time, you must wait for about one minute and then refresh the page to use this feature.
> - When you enable real-time log query, the system creates Log Service projects.

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click Logging and then click **Real-time Log Query**. On the Real-time Log Query tab, click **Activate Now**.

Specify the retention period of logs

By default, logs are retained for seven days. You can modify the retention period based on your business requirements.

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click Logging and then click **Real-time Log Query**. On the Real-time Log Query tab, click **Config Log Retention Time**.
4. In the **Config Log Retention Time** dialog box, modify the retention time. Then, click **OK**.

   Data can be retained for 7 to 3,000 days.

Query real-time logs

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.

- Method 1: Query real-time logs on the **Original Log** tab
  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. Click Logging and then click **Real-time Log Query**. On the Real-time Log Query tab, click the **Original Log** tab and enter the query condition.

  A query condition consists of a query statement and an analytic statement in the `query statement | analytic statement` format. Example: `* | SELECT status, count(*) AS PV GROUP BY status`. The query statement and the analytic statement are separated by a vertical bar (|). The query statement uses proprietary syntax of Log Service.

  A query statement can be individually executed. However, an analytic statement must be executed together with a query statement. In other words, analysis is performed based on the query results or the complete data. For more information, see the "Search and Analysis section of the Log Service" section in *CDS User Guide*.

Statement	Description
Query statement	A query statement specifies one or more query conditions, and then returns the logs that meet the specified conditions. A query statement can be a keyword, a numeric value, a numeric value range, a space character, or an asterisk (*). If you specify a space character or an asterisk (*) as the query statement, no conditions are specified and all logs are returned.
Analytic statement	An analytic statement is used to aggregate or analyze all log data or the log data that meets the specified query conditions in a Logstore.

  iv. Click **15 Minutes(Relative)** to specify the time range for the query statement.

  You can select a relative time, a time frame, or a custom time range. However, the time range that you can specify is only accurate to the minute at most. If you want to use a time range that is accurate to the second, you must specify the time range in the analytic statement. Example: `* | SELECT * FROM log WHERE __time__>1558013658 AND __time__< 1558013660`.
  v. Click **Search & Analyze**.

  Query results contain query and analysis results in a log distribution histogram, on the Raw Logs tab, and on the Graph tab. You can also perform operations on the results. For example, you can configure alerts and create saved searches.

  - Log distribution histogram

    The log distribution histogram shows the distribution of returned logs in different periods of time.

  - Alert

    You can click **Save as Alert** to configure alerts for query and analysis results.

  - Saved search

    You can click **Save Search** to save a query statement as a saved search.

- Raw log

  You can click the Table or Raw Data option on the **Raw Logs** tab to analyze the distribution of a field over a period of time, view the context of the specified log in the raw file, monitor the log content in real time, and extract key log information.

- Graph

  On the **Graph** tab, you can view the visual query and analysis results, add charts to the dashboard, download logs, and configure interactive behaviors.

- LogReduce

  On the **LogReduce** tab, you can click **Enable LogReduce** to cluster similar logs during log collection.

- Method 2: Query real-time logs on the **Dashboard** tab

  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. Click the **Dashboard** tab to analyze logs.

  Dashboard allows you to view four reports that are immediately available.

  - **Access Center**: displays the overall operating status including the PV, UV, traffic, and distribution of access over the Internet.
  - **Audit Center**: displays statistics of object operations including read, write, and delete operations.
  - **Operation Center**: displays statistics of access logs including the number of requests and distribution of failed operations.
  - **Performance Center**: displays statistics of 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.

Disable real-time log query

> **Notice** If you disable real-time log query, the system does not delete Log Service projects. Delete these projects before you disable real-time log query.

1. Log on to the OSS console.
2. Click Buckets. On the Buckets page that appears, click the name of the bucket to which you want to upload objects.
3. Click Logging and then click **Real-time Log Query**. On the Real-time Log Query tab, click **Disable Real-time Log Query**.

4.16. IMG

4.16.1. 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.

Create a style

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 create styles.
3. Click the **Image Processing (IMG)** tab. On the page that appears, click **Create Rule**.
4. In the **Create Rule** panel, configure the style.

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

   - **Basic Edit**: You can use the IMG features by using the graphical user interface (GUI). 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 features that you want to use to process an image. The format is: `image/action1,parame_value1/action2,parame_value2/...`.

     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**.

   After you create a style, you can click **Export** to export the style to a specified local path.

Apply styles

You can perform the following steps to use a created style to process an image object in the current bucket:

1. On the Overview page, click the **Files** tab.
2. Click the name of the image that you want to process.
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 disk.

Simplify IMG URLs that include style parameters

IMG URLs that include style parameters generally include the access URL of the image to process, the style parameter and the name of the style. Example: `https://image-demo.oss-cn-qd-ase-d01-a.mytest-inc.com/example.jpg?x-oss-process=style/small`. You can replace `?x-oss-process=style/` with customized delimiters to simplify the IMG URL. For example, if you specify the delimiter as an exclamation point (!), the IMG URL can be simplified to `https://image-demo.oss-cn-qd-ase-d01-a.mytest-inc.com/example.jpg!small`.

1. In the Buckets page, click the **Image Processing (IMG)** tab.
2. Click **Access Settings**.

3. On the **Access Settings** panel, select one of the **Delimiters**.

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

4. Click **OK**.

4.16.2. Configure source image protection

Object Storage Service (OSS) provides the source image protection feature to protect your images from being used by unauthorized anonymous requesters. After you enable source image protection for your bucket, anonymous requesters can access original images in the bucket only by adding style parameters in the requests or by using signed URLs.

Context

You can use the following methods to access the protected original images:

- Use URLs that contain style parameters in the following format: `https://BucketName.Endpoint/ObjectName?x-oss-process=style/StyleName`.
- Use signed URLs in the following format: `https://BucketName.Endpoint/ObjectName?Signature`.

Procedure

1. Log on to the OSS console.
2. Click the **Image Processing (IMG)** tab. On the page that appears, click **Access Settings**.
3. In the **Access Settings** panel, turn on **Protect Source Image File** and configure the parameters described in the following table.

Parameter	Description
Protected File Extensions	Select a file suffix from the **Protected File Extensions** drop-down list. All objects in the bucket that match the specified suffix are protected.
Delimiters	After you select a custom delimiter, you can use the delimiter to replace ?x-oss-process=style/ to simplify the IMG URL. OSS supports the following delimiters: hyphens (-), underscores (_), forward slashes (/), and exclamation points (!). Click the check box before the delimiter that you want to select. For example, if you set the delimiter to an exclamation point (!), the IMG URL can be simplified to the following format: `http(s)//:BucketName.Endpoint/ObjcetName!StyleName`.

4. Click **OK**.

5.Objects

5.1. Search for objects

You can search for objects whose names contain specific prefixes in buckets or folders in the OSS console.

Prerequisites

Object are uploaded to the bucket. For more information, see Upload objects.

Context

When you search for objects based on a prefix, search strings are case-sensitive and cannot contain forward slashes (/). You can search for objects only in the root folder of the current bucket or in the current folder. Subfolders and objects stored in subfolders cannot be searched.

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 that you want to search for are stored.
3. Click the **Files** tab.
4. Search for objects.
   - Search for objects or folders within the root folder of the bucket

     In the upper-right corner, enter the prefix to search in the search box and press Enter or click the icon to search for related objects. Objects and subfolders whose names contain the specified prefix within the root folder of the bucket are displayed.
   - Search for objects or subfolders within a specified folder

     Click the folder in which the objects or subfolders that you want to search for are stored. In the upper-right corner, enter the prefix to search in the search box and press Enter or click the icon to search for related objects. Objects and subfolders whose names contain the specified prefix within the current folder are displayed.

5.2. Configure object ACLs

You can configure the ACL of an object in the OSS console to control access to the object.

Prerequisites

An object is uploaded to the bucket. For more information, see Upload objects.

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 contains the object whose ACL you want to configure.

3. In the left-side navigation pane, click **Files**.
4. Click the name of the object whose ACL you want to configure. In the **View Details** panel, click **Set ACL** on the right side of **File ACL**.

   You can also choose **More > Set ACL** in the Actions column corresponding to the object whose ACL you want to configure.
5. In the **Set ACL** panel, configure the ACL of the object.

   You can set the ACL of the object to one of the following values:

   - **Inherited from Bucket**: The ACL of the object is the same as that of the bucket.
   - **Private**: Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users cannot access the objects in the bucket without authorization.
   - **Public Read**: Only the owner or authorized users of this bucket can write objects in the bucket. Other users, including anonymous users can only read objects in the bucket.
   - **Public Read/Write**: All users, including anonymous users can read and write objects in the bucket. Fees incurred by such operations are charged to the owner of the bucket. Exercise caution when you set the object ACL to this value.
6. Click **OK**.

5.3. Create folders

You can use the OSS console to create and simulate basic features of folders in Windows. This topic describes how to create a folder by using the OSS console.

Prerequisites

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

Context

OSS does not use a hierarchical structure for objects, but instead uses a flat structure. All elements are stored in buckets as objects. To facilitate object grouping and to simplify management, the OSS console displays objects whose names end with a forward slash (/) as folders. These objects can be uploaded and downloaded. You can use OSS folders in the OSS console in the same manner as you use folders in Windows.

> **Note** The OSS console displays objects whose names end with a forward slash (/) as folders, regardless of whether these objects contain data. The objects can only be downloaded by calling an API operation or by using OSS SDKs.

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 folders.
3. Click the **Files** tab. On the page that appears, click **Create Folder**.
4. In the **Create Folder** panel, enter the folder name.

   The folder name must comply with the following conventions:

- The name can contain only UTF-8 characters and cannot contain emojis.
- The name cannot start with a forward slash (/) or backslash (\). The name cannot contain consecutive forward slashes (/). You can use forward slashes (/) in a folder name to quickly create a subfolder. For example, when you create a folder named *example/test/*, the folder named *example/* is created in the root folder of the bucket and the subfolder named *test/* is created in the example/ folder.
- The name cannot be two consecutive periods (`..`).
- The folder name must be 1 to 254 characters in length.

5. Click **OK**.

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

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

Context

You can configure multiple bucket policies for a bucket. However, the total size of the policies cannot exceed 16 KB.

Method 1: Configure bucket policies by using the GUI

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 bucket policies.
3. On the bucket details page that appears, click **Files**. On the page that appears, click **Authorize**.
4. On the **GUI** tab, click **Authorize**.
5. In the **Authorize** panel, configure the parameters and then click **OK**. The following table describes the parameters that you can configure.

Parameter	Description

Parameter	Description
Applied To	Select the resources that you want to authorize other users to access. ○ **Whole Bucket**: The authorization policy applies to all resources in the whole bucket. ○ **Specified Resource**: The authorization 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 authorize users to access all subdirectories and objects within a directory, add an asterisk (*) at the end of the directory name. For example, to authorize users to access all subdirectories and objects within a directory named abc, enter *abc/**. ■ Configure a bucket policy for a specific object To configure a bucket policy to authorize users to access a specific object, enter the full path of the object that excludes the bucket name. For example, to authorize users to access an object named myphoto.png in the abc directory, enter *abc/myphoto.png*.
Accounts	Select the type of accounts that you want to authorize. ○ **Anonymous Accounts (*)**: Select this option if you want to authorize all users to access the specified resources. ○ **Other Accounts**: Select this option if you want to authorize other Apsara Stack tenant accounts, Resource Access Management (RAM) users, or users that use a temporary token generated by Security Token Service (STS) to access the specified resources. ■ To authorize other Apsara Stack tenant accounts or RAM users to access the specified resources, enter the UIDs of the accounts or RAM users. ■ To authorize users that use a temporary token generated by STS 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 user that uses a temporary token generated by STS is testrole, the UID of the Apsara Stack tenant account that owns 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 authorize all users that use a temporary token generated by STS to access the specified resources, use asterisks (*) as wildcards. For example, enter `arn:sts::*:*/*/*`. **Notice** If an authorized user is a user that uses a temporary token generated by STS, the user cannot access the specified resources on the OSS console.

Parameter	Description
Authorized Operation	You can use the following methods to specify authorized operations: **Basic Settings** and **Advanced Settings**. ○ Basic Settings If you select this method, 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 option. ■ **Read Only**: Authorized users can view, list, and download the specified resources. ■ **Read/Write**: Authorized users can read data from and write data to the specified resources. ■ **Any Operation**: Authorized users can perform all operations on the specified resources. ■ **None**: Authorized users cannot perform operations on the specified resources. **Notice** If multiple bucket policies are configured for a user, the user has all the permissions defined in these policies. However, a policy in which Authorized Operation is set to **None** takes precedence over other policies. For example, if you configure a policy to grant the **Read Only** permission to a user, and then configure another 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**. ○ Advanced Settings If you select this method, you must configure the following parameters: ■ **Effect**: Select Allow or Deny. ■ **Action**: Specify the operation that you want to allow or deny. You can specify any action that is supported by OSS.
Conditions	Optional. You can configure this parameter in both Basic Settings and Advanced Settings to specify the conditions that users must meet to access the specified OSS resources. ○ **Access Method**: Select HTTPS or HTTP. ○ **IP =**: Specify the IP addresses or Classless Inter-Domain Routing (CIDR) blocks that can be used to access the specified OSS resources. Separate multiple IP addresses with commas (,). ○ **IP ≠**: Specify IP addresses or CIDR blocks that cannot be used to access OSS resources. Separate multiple IP addresses with commas (,).

6. Click **OK**.

Method 2: Configure bucket policies by specifying policy syntax

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 bucket policies.
3. On the bucket details page that appears, click **Files**. On the page that appears, click **Authorize**.
4. On the **Syntax** tab, click **Edit**.

   You can specify policy syntax based on your requirements to manage fine-grained permissions. The following examples describe the bucket policies configured by the resource owner whose UID is `174649585760xxxx` in different 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 performing operations on 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/hanghzou/2020/*",
                "acs:oss:*:174649585760xxxx:examplebucket/hangzhou/2015/*"
            ]
        },
        {
            "Action": [
                "oss:ListObjects",
                "oss:ListObjectVersions"
            ],
            "Condition": {
                "StringLike": {
                    "oss:Prefix": [
                        "hanghzou/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 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 and the path of the object. Example: `http://mybucket.oss-cn-beijing.aliyuncs.com/file/myphoto.png`.

- 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 region in which the bucket is located and object path specified in the bucket policy. For more information, see Add OSS paths.

5.5. Delete objects

You can delete uploaded objects in the OSS console when they are no longer needed.

Context

You can delete a single object or batch delete multiple objects. You can batch delete up to 100 objects. To delete specific objects or batch delete more than 100 objects, we recommend that you use API operations or OSS SDKs.

> **Notice** Deleted objects cannot be recovered. Exercise caution when you delete objects.

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 you want to delete are stored.
3. In the left-side navigation pane, click **Files**.
4. Select one or more objects that you want to delete in the object list, and then choose **Batch Operation > Delete**.

   You can also choose **More > Completely Delete** in the Actions column corresponding to the object you want to delete.
5. In the dialog box that appears, click **OK**.

5.6. Manage parts

When you use multipart upload to upload an object, the object is split into several parts. After all of the parts are uploaded to the OSS server, you can call CompleteMultipartUpload to combine the parts into a complete object.

Context

You can also configure lifecycle rules to clear parts that are not needed on a regular basis. For more information, see Manage lifecycle rules.

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 parts you want to delete are stored.
3. Click the **Files** tab. On the page that appears, click **Parts**.

4. In the **Parts** panel, delete the parts.
   - To delete all parts in the bucket, select all parts and then click **Delete All**.
   - To delete specific parts in the bucket, select these parts and then click **Delete**.
5. In the dialog box that appears, click **OK**.

5.7. Configure object tagging

You can configure object tagging to classify objects. Object tagging uses key-value pairs to identify objects. You can perform operations on 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. You can manage multiple objects that have the same tag. For example, you can configure lifecycle rules for objects that have the same tag or authorize Resource Access Management (RAM) users to access objects that have the same tag.

When you configure object tagging, take note of the following items:

- You can add up to 10 tags to an object. The tags added to an object must have unique tag keys.
- A tag key can be up to 128 bytes in length. A tag value can be up to 256 bytes in length.
- Tag keys and tag values are case-sensitive.
- The key and value of a tag can contain letters, digits, spaces, and the following special characters:

  + - = ._:/
- Only the bucket owner and authorized users have read and write permissions on object tags. These permissions are independent of object access control lists (ACLs).
- In cross-region replication (CRR), object tags are also replicated to the destination bucket.

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 object that you want to configure tagging is stored.
3. On the bucket details page, click the **Files** tab.
4. Choose **More** > **Tagging** in the Actions column corresponding to the object to which you want to add tags.
5. In the **Tagging** panel, configure the **Key** and **Value** of the tag.

   You can click **Add** to add up to more 10 tags to the object.
6. Click **OK**.

6.Create single tunnels

You can create single tunnels between OSS and a virtual private cloud (VPC) to access OSS resources from the VPC.

Prerequisites

A VPC and a vSwitch are created.

For more information, see the *Create a VPC* and *Create a vSwitch* topics in *VPC User Guide*.

Procedure

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

Parameter	Required	Description
Organization	Yes	Select the organization of the VPC from which you want to access OSS resources.
Resource Set	Yes	After you select an organization, the resource set is automatically selected based on the organization.
Region	Yes	After you select an organization, a region is automatically selected based on the organization.
Description	No	Enter the description of the single tunnel you want to create. The description cannot exceed 180 characters in length.
VPC	Yes	Select the VPC that you created. You can also click **Create VPC** to create a VPC.
vSwitch	Yes	Select the vSwitch that you created. You can also click **Create vSwitch** to create a vSwitch.

5. Click **Submit**.

7.Add OSS paths

You can add the paths of OSS resources in the console for quicker access.

Prerequisites

A bucket is created. For more information about how to create a bucket, see Create buckets.

Procedure

1. Log on to the OSS console.
2. Click the + icon on the right side of **My OSS Paths**.
3. In the **Add Authorized OSS Path** panel, add a path.

   You can configure the following parameters to add a path.

   - **Region**: Select the region of the bucket in the path that you want to add.
   - **File Path**: Add the path of the resource that you want to access. The path is in the `bucket/object-prefix` format. For example, if the OSS resource that you want to access is the root folder of a bucket named *example*, set File Path to *example*. If the OSS resource that you want to access is the *test* folder in the root folder of the bucket named *example*, set File Path to *example/test/*.