

Alibaba Cloud

Object Storage Service Product Introduction

Document Version: 20220707

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 .
<code>Courier font</code>	Courier font is used for commands	Run the <code>cd /d C:/window</code> command to enter the Windows system folder.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <i>Instance_ID</i>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

Table of Contents

1.What is OSS?	05
2.Features	09
3.Benefits	15
4.Scenarios	18
5.FAQ	20
6.Terms	25
7.Limits	30

1. What is OSS?

Alibaba Cloud Object Storage Service (OSS) is a secure, cost-effective, and highly reliable cloud storage service that allows you to store a large amount of data. OSS is designed to provide 99.999999999% (twelve 9's) data durability and 99.995% data availability. OSS provides multiple storage classes to help you manage and reduce storage costs.

OSS supports RESTful API operations that are independent of the OSS console. You can store and access data from all applications anytime and anywhere.

You can call API operations and use SDKs or OSS migration tools provided by Alibaba Cloud to transfer large amounts of data to and from Alibaba Cloud OSS. You can use OSS buckets of the Standard storage class to store image, audio, and video objects for apps and websites. You can use OSS buckets of the Infrequent Access (IA), Archive, or Cold Archive storage class to store infrequently accessed objects at a low cost and with longer storage period.

Get started with OSS

- Video tutorial

The following video quickly guides you through OSS.

- FAQ

For frequently asked questions about OSS, see [FAQ](#).

- Learning path

[OSS Learning Path](#) helps you get started with OSS. You can learn basic OSS operations, and then perform secondary development by using a variety of API operations, SDKs, and convenient tools.

Terms

- storage class

The storage class of a bucket or an object. OSS provides the following storage classes to cover a variety of data storage scenarios from hot data to cold data: Standard, IA, Archive, and Cold Archive. The Standard storage class provides high-durability, high-availability, and high-performance object storage services that can support frequent data access. The IA storage class is suitable for storing data that is accessed once or twice a month. In addition, the unit price of IA storage is lower than that of Standard storage. The Archive storage class is suitable for storing archive data for at least six months. The Cold Archive storage class is suitable for storing cold data over a long period of time. For more information about storage classes, see [Overview](#).

- bucket

A container used to store objects in OSS. All objects are contained in buckets. You can configure a variety of bucket attributes such as the region, access control list (ACL), and storage class. You can create buckets of different storage classes to store your data.

- object

The basic unit used to store data in OSS. Objects are also known as files. An object consists of object metadata, object data, and an object key. The key of an object uniquely identifies the object in a bucket. Object metadata consists of key-value pairs that describe object attributes, such as the last modified time and the size of the object. You can also specify the user metadata of an object.

- region

The physical location of an OSS data center. You can select a region to create a bucket based on the cost and location of the requests to the bucket. For more information, see [Regions and endpoints](#).

- endpoint

The domain name that is used to access OSS. OSS uses RESTful API operations to provide services. Different regions are accessed by using different endpoints. Access over the internal network or over the Internet within the same region also uses different endpoints. For more information, see [Regions and endpoints](#).

- AccessKey pair

The access credential that is used to authenticate a requester. An AccessKey pair consists of an AccessKey ID and an AccessKey secret. OSS uses symmetric encryption based on AccessKey pairs to verify the identity of a requester. The AccessKey ID is used to identify a user. The AccessKey secret is used to encrypt and verify signature strings. The AccessKey secret must be kept confidential. For more information about how to obtain an AccessKey pair, see [Obtain an AccessKey pair](#).

Common operations

- Create buckets

Before you can upload objects to OSS, you must create a bucket. You can configure various attributes such as the region, ACL, and other metadata for a bucket. For more information about how to create a bucket, see [Create buckets](#).

- Upload objects

After you create a bucket, you can upload objects of different sizes to the bucket. For more information about how to upload objects, see [Upload objects](#).

- Download objects

After you upload objects to a bucket, you can download the objects to the default download path of your browser or the specified local path. For more information about how to download objects, see [Download objects](#).

- List objects

You can list all objects in a bucket. If a large number of objects are stored in a bucket, you can also list some of the objects in the bucket. For more information about how to list objects, see [List objects](#).

- Delete objects

You can manually delete one or more objects that you no longer need. You can also configure lifecycle rules to allow OSS to automatically delete a single object or multiple objects. For more information about how to delete objects, see [Delete objects](#).

Features

- Versioning

OSS allows you to configure versioning for a bucket to protect objects that are stored in the bucket. After you enable versioning for a bucket, data that is overwritten or deleted in the bucket is saved as a previous version. You can use versioning to recover a previous version of an object that is accidentally overwritten or deleted. For more information about versioning, see [Overview](#).

- Bucket Policy

The owner of a bucket can configure bucket policies to grant users different permissions to access the specified OSS resources in the bucket. For example, you can configure bucket policies to authorize other Alibaba Cloud accounts or anonymous users to access or manage all or some resources in your bucket. You can also configure bucket policies to grant full permissions such as read-only, read/write, or full permissions to different RAM users of the same Alibaba Cloud account. This way, different RAM users have different permissions to access or manage resources in your bucket. For more information about how to configure bucket policies, see [Configure bucket policies to authorize other users to access OSS resources](#).

- CRR

Cross-region replication (CRR) provides automatic and asynchronous (near real-time) replication of objects across buckets in different OSS regions. Operations such as creating, overwriting, and deleting objects can be synchronized from a source bucket to a destination bucket. CRR can meet your requirements for cross-region disaster recovery and data replication. For more information about CRR, see [CRR](#).

- Data encryption

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. A header is added in the response to indicate that the object is encrypted on the OSS server. For more information about server-side encryption, see [Server-side encryption](#).

Client-side encryption: Before you upload an object to a bucket, the object is encrypted on the local client. For more information about client-side encryption, see [Client-side encryption](#).

Manage OSS

You can use a variety of methods to upload objects, download objects, and manage OSS.

- Manage OSS by using the OSS console

OSS provides a web-based console. You can log on to the [OSS console](#) to manage your OSS resources. For more information, see [Use Alibaba Cloud accounts to log on to the OSS console](#).

- Manage OSS by using API operations or SDKs

OSS provides RESTful API operations and SDKs for multiple programming languages to facilitate secondary development. For more information, see [List of operations by function](#) and [Overview](#).

- Manage OSS by using tools

OSS provides multiple management tools, such as ossbrowser, ossutil, and ossftp. For more information, see [OSS tools](#).

- Manage OSS by using CSG

OSS uses a flat structure instead of a hierarchical structure to store objects. All elements are stored as objects in buckets. To use OSS in the same manner in which local file systems are used, you can configure Cloud Storage Gateway (CSG). For more information, visit the [CSG product page](#).

Pricing

Traditional storage service providers require you to purchase a reserved amount of storage capacity and network traffic before you use resources. If your storage capacity or traffic usage exceeds the reserved amount, your service is deactivated or you are charged excess fees. At the same time, if your storage capacity or traffic usage does not exceed the reserved amount, you are charged for the full amount of the storage and traffic.

OSS charges you only for the storage capacity and traffic that you consume. Therefore, you do not need to purchase storage capacity and traffic before you use resources. You can take the cost advantages of the flexible infrastructure provided by Alibaba Cloud to grow your business.

For more information about OSS pricing, go to the Pricing tab of the [OSS product page](#). For more information about OSS billing methods, see [Overview](#).

Related services

After you upload your data to OSS, you can use other Alibaba Cloud features and services to manage your data.

The following features and services are frequently used with OSS:

- **Image Processing (IMG):** a service that allows you to perform a variety of operations such as format converting, resizing, cropping, rotating, and adding watermarks to images stored in OSS. For more information, see [IMG implementation modes](#).
- **Elastic Compute Service (ECS):** a cloud computing service that provides elastic and efficient computing capability. For more information, visit the [ECS product page](#).
- **Alibaba Cloud CDN:** a distributed network that caches resources from an origin server to edge nodes in different regions to accelerate content delivery. For more information, visit the [CDN product page](#).
- **E-MapReduce (EMR):** a big data processing solution built on ECS. EMR is developed based on open source Apache Hadoop and Apache Spark to facilitate data analysis and processing. For more information, visit the [E-MapReduce product page](#).
- **ApsaraVideo Media Processing:** a service that converts audio or video objects stored in OSS into files that are suitable for playback on PCs, TVs, and mobile devices. Deep learning results from large amounts of data help ApsaraVideo Media Processing perform multimodal analysis on the audio, text, and images in video files and the content of audio files. ApsaraVideo for Media Processing uses the analysis results to intelligently audit, understand, and edit the content of the audio or video files. For more information, visit the [ApsaraVideo Media Processing product page](#).
- **Data Online Migration:** a service that allows you to migrate data from a third-party storage service such as Amazon Web Services (AWS) and Google Cloud to OSS with ease. For more information, see [Data Online Migration documentation](#).
- **Data Transport:** a service that helps you migrate large amounts of data from local storage to OSS. For example, you can use Data Transport to migrate terabytes or petabytes of local data to OSS when the local network bandwidth is insufficient and expansion costs are high. For more information, see [What is Data Transport?](#)

Alibaba Cloud storage services

In addition to OSS, Alibaba Cloud also provides a variety of storage services such as file storage and block storage for different scenarios. For more information about Alibaba Cloud storage services, see [Overview](#).

For more information about use cases and solutions of Alibaba Cloud storage services, visit the [Alibaba Cloud Storage](#).

2.Features

Object Storage Service (OSS) provides secure, cost-effective, and high-durability services for you to store large amounts of data in the cloud. This topic describes the common scenarios and corresponding features of OSS. You can select the best solution based on your business requirements.

Before you use OSS, we recommend that you have a good understanding of the basic terms used in OSS, including buckets, objects, regions, and endpoints. For more information, see [Terms](#).

The following table describes the features provided by OSS.

Scenario	Feature	References
Upload objects	Before you upload objects to OSS, you must create a bucket in an Alibaba Cloud region to store uploaded objects. After you create a bucket, you can upload objects to the bucket.	<ul style="list-style-type: none">• Create buckets• Create directories• Simple upload
Search for objects	OSS allows you to search for objects and directories to find the object that you want to access in a bucket.	Search for objects
Download objects	After you upload objects to a bucket, you can download the objects to the default download path of your browser or a specified local path.	Download objects
Share objects	After you upload an object to a bucket, you can share the object URL with third parties for downloads or previews.	Share objects
Query objects	OSS provides the data indexing feature to allow you to index the metadata of objects. You can specify the metadata of objects as index conditions to query objects. This way, you can efficiently manage and learn about data structures, perform queries, collect statistics, and manage objects.	Data indexing
Delete objects or directories	OSS allows you to delete one or more objects, directories, and parts at a time. You can delete expired objects at regular intervals to save storage space.	<ul style="list-style-type: none">• Delete objects• Delete directories• Manage parts
Automatically delete multiple objects at a specified period of time	OSS supports lifecycle rules. You can configure lifecycle rules to periodically convert the storage class of non-hot data to Infrequent Access (IA), Archive, or Cold Archive. You can also configure lifecycle rules to delete expired data.	Overview

Scenario	Feature	References
Accelerate the upload and download of data	OSS supports transfer acceleration. Transfer acceleration uses optimal route selection and protocol stack tuning to reduce timeouts in remote transmission and improve user experience.	Transfer acceleration
Recover the data that is accidentally deleted	OSS supports versioning. After you enable versioning for a bucket, OSS saves the overwritten and deleted objects as previous versions in the bucket. You can configure versioning to recover an object to a previous version when the object is accidentally overwritten or deleted.	Overview
Zone-disaster recovery	You can store data in and migrate the data across three zones in the same region by using zone-redundant storage (ZRS). If one zone becomes unavailable, you can continue to access the data that is stored in the other two zones. This feature is designed to provide 99.999999999% (twelve 9's) data durability and 99.995% service availability.	ZRS
Geo-disaster recovery	OSS supports cross-region replication (CRR). You can use CRR to synchronize operations such as creation, update, and deletion of data from the source bucket to the destination bucket in a different region. This way, you can implement geo-disaster recovery.	CRR
Retention policy	OSS supports the Write Once Read Many (WORM) feature. The feature helps prevent objects from being deleted or overwritten within a specified period of time.	合规保留策略

Scenario	Feature	References
Control access to data	<p>OSS supports flexible authorization and authentication mechanisms. You can control access to your OSS resources by using the following methods:</p> <ul style="list-style-type: none"> • ACL: You can set access control lists (ACLs) for buckets and objects, including public read/write, public read, and private. • Bucket policy: You can configure bucket policies to authorize other users to access your OSS resources in the OSS console. For example, you can authorize RAM users of other Alibaba Cloud accounts and anonymous users from specific IP addresses to access your OSS resources. • RAM policy: You can create RAM policies to control access to buckets and directories. OSS provides RAM Policy Editor to generate required RAM policies. For more information, see RAM Policy Editor. • STS temporary authorization: You can use Alibaba Cloud Security Token Service (STS) to grant a third-party application or a RAM user temporary access credentials with a custom validity period. • Hotlink protection: You can configure a Referer whitelist to prevent unauthorized users from accessing your OSS resources. 	<ul style="list-style-type: none"> • Bucket ACL • Configure ACL for objects • Bucket Policy • RAM Policy • Use a temporary credential provided by STS to access OSS • Hotlink protection
Encrypt data	OSS supports client-side and server-side encryption. You can select an encryption method to encrypt and store your data in OSS.	<ul style="list-style-type: none"> • Client-side encryption • Server-side encryption
OSS DDoS protection	OSS DDoS protection is a proxy-based attack mitigation service that integrates OSS with Anti-DDoS Pro and Anti-DDoS Premium. When a bucket with OSS DDoS protection enabled becomes a victim of DDoS attacks, OSS DDoS protection diverts incoming traffic to an Anti-DDoS Pro or Anti-DDoS Premium instance for scrubbing and then redirects normal traffic to the bucket. This ensures the continuity of your business in the event of DDoS attacks.	OSS DDoS protection

Scenario	Feature	References
Integrate with data lake	OSS-HDFS (JindoFS service) is a cloud native data lake storage service. OSS-HDFS is built on unified metadata management capabilities and is fully compatible with Hadoop Distributed File System (HDFS) API operations. The Portable Operating System Interface (POSIX) is supported in OSS-HDFS. This helps OSS-HDFS handle data lake computing scenarios in the big data and AI fields.	Overview of the OSS-HDFS service
Manage data by category	<p>OSS allows you to configure tagging to manage data based on the data category:</p> <ul style="list-style-type: none"> • Configure bucket tags: You can manage data based on bucket tags. For example, you can list buckets that have specified tags and configure the ACLs of buckets that have specified tags. • Configure object tags: You can manage data based on object tags. For example, you can configure lifecycle rules and ACLs for objects that have specified tags. 	<ul style="list-style-type: none"> • Configure bucket tagging • Object tagging
Record the access information for OSS resources	OSS supports logging. You can configure logging to collect access statistics, audit access to OSS, track exceptions, and troubleshoot problems.	<ul style="list-style-type: none"> • Real-time log query • Logging
Use custom domain names to access OSS resources	OSS allows you to map custom domain names to OSS buckets and use the custom domain names to access data in the buckets. If you want to use your custom domain name to access OSS over HTTPS, you can host your Secure Sockets Layer (SSL) certificate in OSS.	<ul style="list-style-type: none"> • Map custom domain names • Host SSL certificates
Static website hosting	OSS supports static website hosting. You can configure static website hosting for your bucket and access static websites by using the bucket domain name.	Overview
CORS	OSS supports cross-origin resource sharing (CORS) in HTML5. CORS allows web application clients that are loaded in one domain name to interact with resources in another domain.	Configure CORS

Scenario	Feature	References
Obtain data from the origin	OSS supports back-to-origin configurations. If your user accesses data in a bucket that has no back-to-origin rules configured and the data does not exist, 404 Not Found is returned. However, if you configure back-to-origin rules that contain the correct origin URLs, users can obtain the data from the origin based on the back-to-origin rules. You can configure back-to-origin rules for hot data migration and specific request redirection.	Overview
Query object metadata	OSS supports the inventory feature. You can configure inventories for buckets to export the metadata of specified objects, including the object sizes and encryption status.	Bucket inventory
Modify HTTP headers	OSS allows you to modify the HTTP headers of objects. You can configure HTTP headers to specify custom HTTP request policies, such as the cache policy and forced object download policy.	Configure object metadata
View the details of resource usage	OSS supports the monitoring feature. You can use the monitoring feature to view real-time information about OSS service usage, such as the running status and performance of the system.	Overview
Traffic throttling	OSS supports single-connection bandwidth throttling. You can configure single-connection bandwidth throttling for upload, download, and copy operations on OSS objects to ensure sufficient bandwidth for your other applications.	Single-connection bandwidth throttling
Analyze and process data	OSS provides Image Processing (IMG) and video snapshot capturing for you to analyze and process data stored in OSS: <ul style="list-style-type: none"> • IMG: You can perform operations on images stored in OSS, such as format conversion, cropping, scaling, rotating, and watermarking. You can also encapsulate multiple operations into a style to process images stored in OSS. • Video snapshot capturing: You can capture images from video objects in the H.264 format. 	<ul style="list-style-type: none"> • Overview • Video snapshots
Use tools to manage OSS resources	OSS provides graphical, command-line, file mounting, and FTP tools for you to manage your OSS resources.	OSS tools

Scenario	Feature	References
Use OSS SDKs to manage OSS resources	OSS provides SDKs for a variety of programming languages to facilitate custom development.	Overview

3.Benefits

Object Storage Service (OSS) is a secure, cost-effective, and high-durability service that enables you to store large amounts of data in Alibaba Cloud. This topic compares OSS with the traditional self-managed server storage to help you better understand the benefits of OSS.

Advantages of OSS over self-managed server storage

Item	OSS	Self-managed server storage
Durability	<p>OSS is the core infrastructure of data storage for Alibaba Group. This reliable and highly available service is one of the backbone services that ensure stability during the peak hours of Double 11. OSS comes with a multi-redundant architecture to provide reliable data storage. In addition, OSS is built on a high-availability architecture to eliminate single points of failure (SPOFs) and ensure the continuity of your services.</p> <ul style="list-style-type: none">• Provides service availability of at least 99.995%.• Provides data durability of at least 99.9999999999% (twelve 9's).• Automatically expands capacities without affecting your services.• Offers automatic data redundancy.	<ul style="list-style-type: none">• Prone to errors due to hardware failures. If a disk has a bad sector, data may be lost.• Requires manual restoration of data, which can be a complex, time-consuming, and labor-intensive process.
Security	<ul style="list-style-type: none">• Provides multi-level security protection for enterprises, including features such as server-side encryption, client-side encryption, hotlink protection, bucket policies that specify IP address blacklists or whitelists, fine-grained access control, log audit, and the Write Once Read Many (WORM) strategy.• Provides resource isolation mechanisms for multiple tenants and supports geo-disaster recovery.• Complies with the regulations of multiple organizations including the U.S. Securities and Exchange Commission (SEC) and Financial Industry Regulatory Authority, Inc. (FINRA), which delivers services that can meet the requirements of your enterprise for data security and compliance.	<ul style="list-style-type: none">• Requires additional scrubbing devices and blackhole policy-related services.• Requires a separate security mechanism.

Item	OSS	Self-managed server storage
Cost	<ul style="list-style-type: none"> Allows you to use multi-line Border Gateway Protocol (BGP) to connect to operator backbone network lines with sufficient bandwidth resources. Upstream traffic is free of charge. Requires no hosting fees or the need to hire O&M personnel. 	<ul style="list-style-type: none"> Requires manual expansion of storage space due to limited capacity of hard disks. Offers limited bandwidth due to slow single-line or double-line access. Requires the manual expansion during peak hours. Requires dedicated O&M personnel, which results in high costs.
Intelligent storage	Provides multiple data processing capabilities, such as Image Processing (IMG), video snapshot, document preview, image scenario recognition, facial recognition, and SQL in-place query. OSS can be seamlessly integrated with the Hadoop ecosystem and Alibaba Cloud services such as Function Compute, E-MapReduce (EMR), Data Lake Analytics (DLA), BatchCompute, MaxCompute, and Database Backup (DBS) to manage and analyze data of enterprises.	Requires separate purchase and deployment of data processing capabilities.

Additional benefits of OSS

- Ease of use
 - OSS provides standard RESTful API operations, a wide range of SDKs, client tools, and the OSS console. You can upload, download, retrieve, and manage large amounts of data used for websites and mobile apps in the same way you use regular file systems.
 - The capacity of each bucket is unlimited. Therefore, you can expand your buckets in OSS based on your requirements.
 - OSS supports streaming writes and reads. It is suitable for business scenarios that require simultaneous write and read of large files such as videos.
 - OSS supports lifecycle management. You can configure lifecycle rules to batch delete expired objects or convert the storage classes of expired objects to cost-effective Infrequent Access (IA), Archive, or Cold Archive.
- Powerful and flexible security mechanisms
 - OSS provides Security Token Service (STS) and URL authentication and authorization. OSS also supports IP address blacklists or whitelists, hotlink protection, and Resource Access Management (RAM).
 - OSS provides resource isolation mechanisms for users. You can also use the multi-cluster synchronization service.
 - OSS provides server-side encryption, client-side encryption, and encrypted transmission based on Secure Sockets Layer (SSL) or Transport Layer Security (TLS) to protect data from potential security risks on the cloud.

- OSS provides the versioning feature to prevent objects from being accidentally deleted or overwritten.

- Data redundancy mechanism

OSS uses a data redundant storage mechanism to store copies of each object in multiple devices of different facilities in the same region. This way, OSS ensures data durability and availability even if hardware failures occur.

- Operations performed on objects in OSS are strongly consistent. For example, after you upload or copy an object, the uploaded or copied object can be immediately read and copies of the object are written to multiple devices for redundancy.
- To ensure that data is completely transmitted, OSS calculates the checksum of the network traffic packets to check for errors when packets are transmitted between the client and the server.
- The data redundancy mechanism of OSS can prevent data loss even when two storage facilities are damaged at the same time.
 - After data is stored in OSS, OSS regularly checks whether copies of the data are lost. Then, OSS recovers lost copies to ensure data durability and data availability.
 - OSS periodically verifies the integrity of data to detect data corruption caused by errors such as hardware failures. If data is partially corrupted or lost, OSS uses the other copies to reconstruct and repair the corrupted data.

- Rich and powerful value-added services

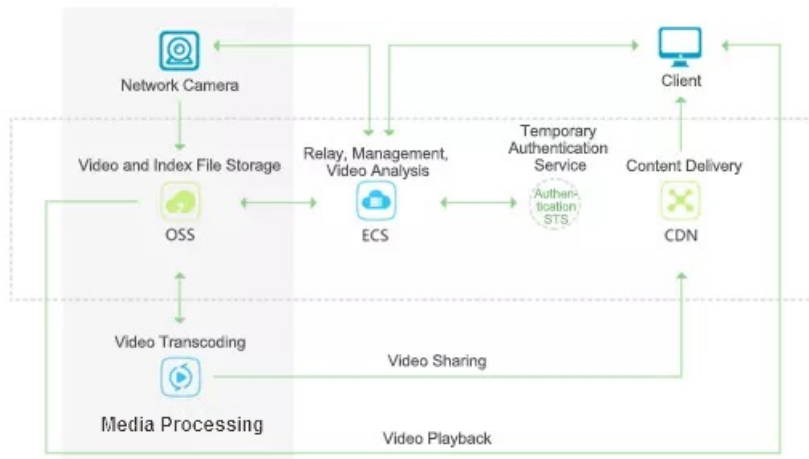
- IMG: supports format conversion, thumbnails, cropping, watermarking, and resizing for objects in formats such as JPG, PNG, BMP, GIF, WebP, and TIFF.
- Audio or video transcoding: provides high-quality, high-speed, and parallel audio or video transcoding capabilities. This way, your audio or video files can be played on different terminal devices.
- Transfer acceleration over the Internet: provides the transfer acceleration service, which uses optimal route selection and protocol stack tuning to reduce timeouts in remote transmission and improve user experience. For more information, see [Transfer acceleration](#).
- Accelerated content delivery: uses OSS as the origin and works with Content Delivery Network (CDN) to improve user experience when the same object is repeatedly downloaded.

4.Scenarios

This topic describes the application scenarios of OSS.

Massive storage for image, audio, and video applications

OSS can be used to store large amounts of data, such as images, audio and video, and logs. You can use a variety of devices, websites, and mobile applications to read data from or write data to OSS. You can write data to OSS by uploading files or using streams.




Dynamic and static content separation for websites and mobile applications

By using the high bandwidth provided by OSS, you can concurrently download a large amount of data from OSS over the Internet. OSS provides **Transfer acceleration** that accelerates uploads and downloads across countries and continents and improves user experience. You can also use OSS with Content Delivery Network (CDN) to deliver static data stored in OSS to CDN edge nodes. Users can access the data cached on CDN edge nodes instead of the same data stored in OSS. This way, the concurrent and repeated download of an object within the same region can be accelerated.

□

Cloud data processing

After you upload objects to OSS, you can use ApsaraVideo for Media Processing and Image Processing (IMG) to process the uploaded data.



5.FAQ

This topic provides answers to frequently asked questions about Object Storage Service (OSS).

General FAQ

- What is Alibaba Cloud OSS?

Alibaba Cloud OSS is a secure, cost-effective, and scalable storage service that has high durability and is capable of processing large amounts of data. OSS is designed for data durability of at least 99.999999999% (twelve 9's) and provides service availability of at least 99.995%.

- What are the features of OSS?

OSS supports RESTful API operations that are independent of the OSS console. You can store and access data from all applications anytime and anywhere. OSS is highly scalable. You are charged only for the resources that you use. You can scale OSS resources based on your requirements without compromising performance or durability.

You can call API operations and use SDKs or OSS migration tools provided by Alibaba Cloud to transfer large amounts of data to and from Alibaba Cloud OSS. You can use OSS buckets of the Standard storage class to store image, audio, and video objects for apps and websites. You can use OSS buckets of the Infrequent Access (IA), Archive, or Cold Archive storage class to store infrequently accessed objects at a low cost and with longer storage period.

For more information about the features of OSS, see [Features](#).

- Who are the intended users of OSS?

OSS is an ideal solution for users who need to store large amounts of data. These users include app and software developers, game development enterprises, and webmasters of online communities, media-sharing sites, and e-commerce websites.

- What data does OSS store?

OSS is suitable for storing attachments, high-definition images, audio and video objects, and backup objects for forums, websites, and software applications, as well as objects for a variety of applications, file synchronization software, and online storage systems.

- What are the advantages of OSS over local storage solutions?

OSS allows developers to make full use of the economy of scale of Alibaba Cloud at minimal cost without additional investments or affecting performance. Developers can focus on their own innovations without performance bottlenecks and security risks that come with increasing business scales. OSS is cost-effective and easy to use.

- What is the upper limit of the data volume that OSS can support?

OSS does not impose limits on the total storage capacity and the capacity of 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 in size, use multipart upload, resumable upload, ossbrowser, or ossutil. For more information, see [Multipart upload](#), [Quick start](#), and [Overview](#).

- What storage classes does OSS provide?

The storage class of a bucket or an object. OSS provides the following storage classes to cover a variety of data storage scenarios from hot data to cold data: Standard, IA, Archive, and Cold Archive. The Standard storage class provides high-durability, high-availability, and high-performance object storage services that can support frequent data access. The IA storage class is suitable for storing data that is accessed once or twice a month. In addition, the unit price of IA storage is lower than that of Standard storage. The Archive storage class is suitable for storing archive data for at least six months. The Cold Archive storage class is suitable for storing cold data over a long period of time. For more information about storage classes, see [Overview](#).

- Will Alibaba Cloud use the data that I store in OSS?

Alibaba Cloud does not use or release your data without authorization. Alibaba Cloud processes user data only based on your service requirements or requirements of laws and regulations. For more information, see [Service terms](#).

- Does Alibaba Cloud use OSS to store their own data?

Yes, Alibaba Cloud developers also use OSS to store authorized data for many projects. These projects rely on OSS to perform key business operations.

- How does OSS ensure availability when my OSS service experiences bursts of traffic from applications?

OSS is designed to address bursts of traffic from Internet applications. Pay-as-you-go is used as a billing method. To ensure service continuity during bursts of traffic, OSS does not impose capacity limits. OSS ensures load balancing so that applications are not affected by bursts of traffic.

- How is OSS data organized?

OSS is a distributed object storage service that stores data as key-value pairs. When you store an object, you must specify an object name (key). The key can then be used to obtain the content of the object.

Keys can also be used to simulate certain features of directories. OSS uses a flat structure for objects instead of a hierarchical structure. All elements are stored as objects in buckets. However, OSS supports directories as a concept to group objects and simplify management. When you use APIs or SDKs to configure an object, you can specify the key value, a full name that includes a prefix for the object to manage other objects. For example, if you set the key of an object to `dir/example.jpg`, a directory named `dir` is created in the current bucket and an object named `example.jpg` is created in the directory. If you delete the `dir/example.jpg` object, the `dir` directory is also deleted.

- How do I get started with OSS?

Before you use OSS, ensure that you have created an Alibaba Cloud account and completed real-name verification. If you do not have an Alibaba Cloud account, the system prompts you to create an Alibaba Cloud account when you activate OSS. For more information, see [Create Your Alibaba Cloud Account](#). After you create the Alibaba Cloud account, visit the [Object Storage Service product homepage](#). Click **Activate Now**. By default, the billing method after you activate OSS is Pay-as-you-go. To further minimize OSS costs, we recommend that you purchase resource plans. For more information, see [Buy OSS resource plans \(subscription\)](#).

You can get started with OSS by using the OSS console, graphical management tool, command-line management tool, and SDKs for various programming languages. For more information, see [Get started with OSS](#).

FAQ about Alibaba Cloud regions

- What is an Alibaba Cloud region?

An Alibaba Cloud region is a geographical region that contains multiple geographically isolated zones. These zones are connected to each other by using networks that feature short latency, high throughput, and high redundancy.

- What is a zone?

Each region contains several isolated locations known as zones. The power supply and network of each zone are independent. The network latency between instances within the same zone is shorter. Zones within the same region can access each other. However, faults that occur in one zone do not affect other zones.

- How do I determine which region to store my data?

When you select a region, we recommend that you consider factors such as physical locations, relationships between cloud services, and resource prices. For more information, see [Choose an OSS region](#).

FAQ about billing

- How does OSS charge fees?

OSS provides the pay-as-you-go billing method that allows you to pay for resources you use. No minimum usage fees are imposed. You can also purchase resource plans. Resource plans are used to deduct fees incurred by resource usage. In most cases, resource plans are more cost-effective. For more information about prices, visit the [OSS pricing page](#).

- How does OSS charge fees when other accounts are used to access my OSS resource?

When other accounts access your OSS resources, OSS charges you based on standard pricing. You can enable the pay-by-requester mode for your bucket so that requesters pay the fees generated by sending requests and downloading OSS data. For more information, see [Enable pay-by-requester](#).

FAQ about data security and protection

- Is data stored in OSS secure?

OSS is secure. By default, only the resource owner can access resources within a bucket. OSS provides user identity verification to control access to data. You can use various access control policies on the bucket or object level, such as access control lists (ACLs), to grant specific permissions to specific users and user groups. The OSS console displays buckets that are available to public access. You can set the bucket ACL to private if you do not want other users to access your bucket or object. If you set the bucket ACL to public read or public read/write, a message that indicates security risks appears. For more information about the security of OSS, see [Overview](#).

- How do I control access to my OSS data?

OSS provides multiple access control methods, including ACLs, RAM policies, and bucket policies, for access to objects stored in buckets. For more information, see [Overview](#).

- What data encryption methods does OSS provide?

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. A header is added in the response to indicate that the object is encrypted on the OSS server. For more information about server-side encryption, see [Server-side encryption](#).

Client-side encryption: Before you upload an object to a bucket, the object is encrypted on the local client. For more information about client-side encryption, see [Client-side encryption](#).

- How do I prevent data stored in buckets from being accidentally deleted or overwritten?

OSS allows you to configure versioning for a bucket to protect objects that are stored in the bucket. After you enable versioning for a bucket, data that is overwritten or deleted in the bucket is saved as a previous version. You can use versioning to recover a previous version of an object that is accidentally overwritten or deleted. For more information about versioning, see [Overview](#).

- What is a retention policy?

OSS supports the Write Once Read Many (WORM) strategy that prevents an object from being deleted or overwritten for a specified period of time. You can configure time-based retention policies for buckets. After a retention policy is configured and locked for a bucket, you can read objects from or upload objects to the bucket. However, objects in the bucket or the retention policy cannot be deleted within the retention period of the retention policy. You can delete the objects only after the retention period expires.

You can configure retention policies when you need to store infrequently accessed important data that is not allowed to be modified or deleted. Such data includes medical records, technical documents, and contracts. You can store these objects in a specified bucket and configure a retention policy for the bucket.

FAQ about data replication

- How do I replicate data from a bucket to another bucket in a different region?

You can configure multiple cross-region replication (CRR) rules for a bucket to store multiple copies of the data in different regions. CRR provides automatic and asynchronous (near real-time) replication of objects across buckets in different OSS regions. Operations such as creating, overwriting, and deleting objects can be synchronized from a source bucket to a destination bucket.

- What are the advantages of CRR?

- Compliance requirements: Although OSS stores multiple replicas of each object in physical disks, replicas must be stored at a distance from each other to comply with regulations. CRR allows you to replicate data between geographically distant OSS data centers to satisfy these compliance requirements.
- Minimum latency: You have users who are located in two geographical locations. To minimize the latency when the users access objects, you can maintain replicas of objects in OSS data centers that are geographically closer to these users.
- Data backup and disaster recovery: You have high requirements for data security and availability, and want to explicitly maintain replicas of all written data in a second data center. If one OSS data center is damaged in a catastrophic event such as an earthquake or a tsunami, you can use backup data from the other data center.
- Data replication: For business reasons, you may need to migrate data from one OSS data center to another data center.
- Operational reasons: You have compute clusters deployed in two different data centers that need to analyze the same group of objects. You can choose to maintain object replicas in these regions.

- How does OSS charge fees when I use CRR?

After CRR is enabled, cross-region traffic is generated when you replicate objects across buckets in the source and destination regions. You are charged for the traffic that is generated when you use CRR. Each time an object is synchronized, OSS calculates the number of requests and charges the requests on a pay-as-you-go basis. The traffic that is generated when you use CRR can be charged only on a pay-as-you-go basis. Resource plans are unavailable for CRR.

FAQ about data query

How do I query data in OSS?

OSS supports the `SelectObject` operation that allows you to use SQL statements to query specific data in a CSV or JSON object instead of querying the entire object. The `SelectObject` operation simplifies the process used to query data and filter the data into smaller and more specific data sets. Therefore, this operation is suitable for the multipart query of large objects, query of JSON objects, and analysis of log objects. For more information about data query, see [Call the `SelectObject` operation to query objects](#).

FAQ about storage management

- What is OSS lifecycle management? How do I use lifecycle management to minimize OSS storage costs?

You can configure a lifecycle rule to regularly delete objects that are no longer accessed or convert the storage class of cold data to IA, Archive, or Cold Archive. This makes data management easier and saves storage costs. For example, you can configure lifecycle rules in the following scenarios:

- A medical institution stores its medical records in OSS. These objects are occasionally accessed within six months after they are uploaded, and almost never after that. In this case, you can configure a lifecycle rule to convert the storage class of these objects to Archive 180 days after they are uploaded.
- A company stores the call records of its service hotline in OSS. These objects are frequently accessed within the first two months, occasionally after two months, and almost never after six months. After two years, these objects no longer need to be stored. In this case, you can configure a lifecycle rule to convert the storage class of these objects to IA 60 days after they are uploaded and to Archive 180 days after they are uploaded, and delete them 730 days after they are uploaded.
- You can manually delete up to 1,000 objects each time. If a bucket contains more than 1,000 objects, you must delete the objects multiple times. In this case, you can configure a lifecycle rule that is used to delete all objects in the bucket the next day. This way, all objects in the bucket can be deleted the next day.

For more information about lifecycle management, see [Lifecycle rules based on the last modified time](#).

- How do I regularly obtain the information about objects stored in a bucket?

You can use the bucket inventory feature to export the information about specific objects in a bucket on a daily or weekly basis. Exported object information includes the number, sizes, storage classes, and encryption status of the objects. For more information about the bucket inventory feature, see [Bucket inventory](#).

6. Terms

This topic describes several basic terms used in Object Storage Service (OSS).

bucket

The container for OSS objects. Each object in OSS is contained in a bucket. You can configure various attributes of a bucket, including the region, access control list (ACL), and storage class. You can create buckets of different storage classes to store data based on your requirements.

- OSS does not use a hierarchical structure for objects, but instead uses a flat structure. Each object belongs to a bucket.
- You can create multiple buckets.
- A bucket name must be unique in OSS within an Alibaba Cloud account. Bucket names cannot be changed after the buckets are created.
- A bucket can contain an infinite number of objects.

OSS supports the following bucket naming conventions:

- The name can contain only lowercase letters, digits, and hyphens (-).
- The name must start and end with a lowercase letter or a digit.
- The name must be 3 to 63 characters in length.


object

The basic unit for data operations in OSS. Objects are also known as files. OSS does not use a hierarchical structure for objects, but instead uses a flat structure. All elements are stored as objects in buckets. However, OSS supports directories as a concept to group objects and simplify management. An object is composed of object metadata, user data, and a key. A key is used to identify an object in a bucket. Object metadata is a group of key-value pairs that define the properties of an object, such as the last modified time and the object size. You can also assign user metadata to the object.

The lifecycle of an object starts when the object is uploaded, and ends when the object is deleted. Throughout the lifecycle, content can be appended only to objects created by using append upload. If you want to modify the content of an object, you must upload a new object that has the same name as the existing object to replace the existing object.

The name of an object must comply with the following conventions:

- The name must be encoded in UTF-8.
- The name must be 1 to 1,023 characters in length.
- The name cannot start with a forward slash (/) or a backslash (\).

 **Note** Object names are case-sensitive. Unless otherwise stated, objects and files mentioned in OSS documents are called objects.

ObjectKey

In SDKs for different programming languages, ObjectKey, Key, and ObjectName indicate the full path of the object. You must specify the full path of an object when you perform operations on the object. For example, when you upload an object to a bucket, ObjectKey indicates the full path that includes the extension of the object. For example, you can set ObjectKey to abc/efg/123.jpg.

region

The physical location of OSS resources. When you create a bucket, you can select a region based on the cost and source of the requests. In most cases, the closer a user is located from an OSS region, the faster the access. For more information, see [Regions and endpoints](#).

A region is specified when a bucket is created. After a bucket is created, its region cannot be changed. All objects in this bucket are stored in the corresponding region. Regions are configured for buckets instead of objects.

endpoint

The domain name that is used to access OSS. OSS uses HTTP RESTful APIs to provide services. Different regions are accessed by using different endpoints. A region has different endpoints for access over the internal network and for access over the Internet. For example, the public endpoint used to access OSS data in the China (Hangzhou) region is `oss-cn-hangzhou.aliyuncs.com`, and the internal endpoint is `oss-cn-hangzhou-internal.aliyuncs.com`. For more information, see [Regions and endpoints](#).

AccessKey pair

The access credential that is used to identify the requester. An AccessKey pair consists of an AccessKey ID and an AccessKey secret. OSS uses an AccessKey pair to implement symmetric encryption and verify the identity of a requester. The AccessKey ID is used to identify a user. The AccessKey secret is used to encrypt and verify the signature string. The AccessKey secret must be kept confidential. OSS supports AccessKey pairs obtained by using the following methods:

- The bucket owner applies for an AccessKey pair.
- The bucket owner uses Resource Access Management (RAM) to assign an AccessKey pair to a third party.
- The bucket owner uses Security Token Service (STS) to assign an AccessKey pair to a third party.

For more information about AccessKey pairs, see [Obtain an AccessKey pair](#).

strong consistency

The feature that requires object operations in OSS be atomic, which indicates that operations can either succeed or fail without intermediate states. To ensure that users can access only complete objects, OSS does not return partially uploaded objects.

Object operations in OSS are highly consistent. For example, when a user receives an upload (PUT) success response, the uploaded object can be immediately read, and copies of the object are created for redundancy. Therefore, the situations where data is not obtained when a user performs a read-after-write operation do not exist. The same is true for delete operations. After a user deletes an object, the object and its copies no longer exist.

data redundancy mechanism

The data redundancy mechanism that is implemented based on erasure coding and multiple replicas. Copies of each object are stored in different servers within the same region. This way, data durability and availability are ensured when hardware failures occur.

- Operations on objects in OSS are highly consistent. For example, when a user receives an upload or copy success response, the uploaded object can be immediately read, and copies of the object are created for redundancy.
- To ensure complete data transmission, OSS calculates the checksum of the network traffic packets and checks for errors when packets are transmitted between the client and the server.

- The data redundancy mechanism of OSS can prevent data loss even when two storage facilities are damaged at the same time.
 - After data is stored in OSS, OSS regularly checks whether copies of the data are lost. Then, OSS recovers lost copies to ensure data durability and data availability.
 - OSS periodically verifies the integrity of data to detect data corruption caused by errors and hardware failures. If data is partially corrupted or lost, OSS recovers the data by using copies of the data.

Comparison between OSS and file systems

Item	OSS	File system
Data model	OSS is a distributed object storage service that stores data as key-value pairs.	A file system uses a typical tree structure for directory indexing.
Data retrieval	<p>Objects are retrieved based on unique object names (keys).</p> <p>For example, the object name test1/test.jpg does not necessarily indicate that the object is stored in a directory named test1. In OSS, test1/test.jpg is only a string. test1/test.jpg is not essentially different from example.jpg. Therefore, similar amounts of resources are consumed regardless of which object you access.</p>	To access a file named test1/test.jpg, you must first access the test1 directory and then query the test.jpg file in this directory.
Advantages	OSS supports concurrent access from a large number of users.	The file system supports modifications on files, such as modifying the content at a specified offset location or truncating the end of a file. It also supports directory operations such as rename, delete, and move directories.

Item	OSS	File system
Disadvantages	<p>Objects stored in OSS cannot be modified. A specific operation must be called to append an object. Generated objects are different from objects uploaded by using other methods. To modify an object, you must upload the entire object again.</p> <p>OSS can simulate features similar to those of directories. However, such operations consume a large amount of resources. For example, if you want to rename the test1 directory to test2, OSS must copy all objects whose names start with test1/. Then, OSS creates objects whose names start with test2/. This operation consumes a large amount of resources. Therefore, we recommend that you do not perform such operations in OSS.</p>	<p>The performance of the file system is subject to the performance of a single device. More files and directories in the file system consume greater resources and take longer processing time.</p>

We recommend that you do not map operations on OSS objects to file systems because it is inefficient. If you mount OSS as a file system, we recommend that you perform only add, delete, and read operations on objects. You can take advantage of OSS to process and store large amounts of unstructured data such as images, videos, and documents.

The following table describes the differences in some terms between OSS and file systems.


OSS	File system
Object	file
Bucket	home directory
Region	N/A
Endpoint	N/A
AccessKey	N/A
N/A	multilevel directory
GetService	obtain the list of home directories


OSS	File system
GetBucket	obtain the list of files
PutObject	add a file
AppendObject	append data to an existing file
GetObject	read a file
DeleteObject	delete a file
N/A	modify file content
CopyObject (same destination and source objects)	modify file attributes
CopyObject (different destination and source objects)	copy a file
N/A	rename a file

7.Limits

This topic describes the performance metrics and limits of Object Storage Service (OSS).

The following table describes the performance metrics and limits of OSS.

Item	Description
Bandwidth	<p>The default bandwidth limit is 10 Gbit/s for regions inside the Chinese mainland and 5 Gbit/s for regions outside the Chinese mainland. If this limit is reached, requests are throttled.</p> <div> <p> Note When a request is throttled, the response to the request contains the <code>x-oss-qos-delay-time: number</code> header in which <code>number</code> indicates the duration over which the request is throttled. Unit: ms. For upload requests, the exact duration over which a request is throttled is returned. For download requests, the estimated duration over which a request is throttled is returned. The duration is estimated based on the extent of throttling and the size of the downloaded objects.</p> </div> <p>If you require a higher bandwidth (10 Gbit/s to 100 Gbit/s) for your business such as offline big data processing, contact technical support.</p>
Queries per second (QPS)	<p>The total QPS limit for a single account is 10,000. The actual values are different in different read and write modes:</p> <ul style="list-style-type: none"> Sequential read and write: 2,000 <p>If you upload a large number of objects with sequential prefixes such as timestamps and letters in the object names, many object indexes may be stored in a single partition. In this case, if you send a large number of requests to query these objects, the latency may increase. We recommend that you do not upload a large number of objects that contain sequential prefixes in their names. For more information about how to change sequential prefixes to random prefixes, see OSS performance and scalability best practices.</p> <ul style="list-style-type: none"> Non-sequential read and write: 10,000 <p>If you require a higher QPS, contact technical support.</p>
Bucket	<ul style="list-style-type: none"> You can use an Alibaba Cloud account to create up to 100 buckets in the same region. A bucket name must be globally unique within OSS. For more information about the naming conventions of buckets, see Bucket naming conventions. After a bucket is created, its name, region, storage class, and redundancy type cannot be modified. OSS does not impose limits on the capacity of a bucket.

Item	Description
Object	<ul style="list-style-type: none"> The size of the object that you want to upload When you use simple upload, form upload, or append upload to upload a single object, the object cannot exceed 5 GB in size. When you use multipart upload to upload a single object, the object cannot exceed 48.8 TB in size. The size of the object that you want to rename You can rename only objects that are up to 1 GB in size in the OSS console. To rename objects that exceed 1 GB in size, we recommend that you use ossutil. The number of objects that you want to delete You can delete up to 100 objects at a time in the OSS console. You can delete up to 1,000 objects at a time when you use OSS SDKs, ossbrowser, or ossutil. <div>  Warning Deleted objects cannot be recovered. Exercise caution when you delete objects. </div> <ul style="list-style-type: none"> Whether existing objects are overwritten by objects with the same names By default, when you upload an object that has the same name as an existing object, the existing object is overwritten. To prevent objects from being accidentally overwritten, you can enable versioning for the bucket in which the objects are stored. You can also include the x-oss-forbid-overwrite header in upload requests and set the value of the header to <i>true</i>.
Data restoration	<p>You must restore an Archive object or a Cold Archive object before you can access the object.</p> <ul style="list-style-type: none"> It takes about a minute to restore an Archive object. For a Cold Archive object, the restoration period varies based on the restoration priority of the object. <ul style="list-style-type: none"> Expedited: The object is restored within one hour. Standard: The object is restored within 2 to 5 hours. Bulk: The object is restored within 5 to 12 hours.
Custom domain name	<ul style="list-style-type: none"> For a domain name that is mapped to buckets located in regions inside the Chinese mainland, you must apply for an ICP filing at the Ministry of Industry and Information Technology (MIIT). Up to 100 domain names can be mapped to a bucket. However, a domain name can be mapped only to a single bucket. OSS does not impose limits on the number of domain names that can be mapped to the buckets of an Alibaba Cloud account.
Lifecycle rules	You can configure up to 1,000 lifecycle rules for a bucket.

Item	Description
Back-to-origin rules	<ul style="list-style-type: none"> You can configure up to 20 back-to-origin rules for a bucket. In regions inside the Chinese mainland and the China (Hong Kong) region, the default QPS for mirroring-based back-to-origin requests is 2,000, and the default bandwidth is 2 Gbit/s. In regions outside China, the default QPS for mirroring-based back-to-origin requests is 1,000, and the default bandwidth is 1 Gbit/s.
Image Processing (IMG)	<ul style="list-style-type: none"> Limits on images <ul style="list-style-type: none"> Source images <ul style="list-style-type: none"> Only JPG, PNG, BMP, GIF, WebP, and TIFF images are supported. The size of a source image cannot exceed 20 MB. For the rotate operation on images, the height or width of the source image cannot exceed 4,096 pixels. For other operations on images, the width or height of the source image cannot exceed 30,000 pixels, and the total number of pixels of the source image cannot exceed 250 million. The total number of pixels of a dynamic image, such as a GIF image, is calculated by using the following formula: $\text{Width} \times \text{Height} \times \text{Number of image frames}$. The total number of pixels of a static image, such as a PNG image, is calculated by using the following formula: $\text{Width} \times \text{Height}$. Resized images <p>The width or height of a resized image cannot exceed 16,384 pixels. The total number of pixels of the resized image cannot exceed 16,777,216.</p> Limits on image styles <p>You can create up to 50 image styles in each bucket. To create more styles for a bucket, contact technical support.</p>

Item	Description
Resource plans	<ul style="list-style-type: none">• Resource plans that you purchase for specific regions can be used to offset the fees generated by the resource usage only within the specified regions.• You can use resource plans to offset the fees generated by the resource usage in the regions that you specify when you purchase the resource plans.• You can purchase only one storage plan for each validity period. However, you can upgrade the storage plan that you purchase. For more information about how to upgrade a storage plan, see Overview.• You can purchase multiple transfer acceleration plans and back-to-origin traffic plans within the same validity period to offset your fees. However, these plans cannot be upgraded or renewed.• You can purchase multiple downstream traffic plans within the same validity period to offset your fees. These plans can be renewed but cannot be upgraded.• No resource plans are available to offset API operation calling fees, data processing fees, and cross-region replication (CRR) traffic fees. You are charged these fees on a pay-as-you-go basis. For more information, see API fees, Data processing fees, and Traffic fees.