

Alibaba Cloud

Object Storage Service Product Introduction

Document Version: 20201104

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 <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.Benefits	08
3.Scenarios	12
4.Features	13
5.Terms	17
6.Limits	22

1. What is OSS?

Object Storage Service (OSS) is a secure, cost-effective, and highly reliable cloud storage service provided by Alibaba Cloud. It enables you to store a large amount of data in the cloud. OSS is designed to provide 99.999999999 (twelve 9s) durability and 99.995% availability of data over a given year.

OSS supports RESTful API operations that are independent of the console. You can store and access any type of data anytime, anywhere, and from any application.

You can use API operations and SDKs provided by Alibaba Cloud or OSS migration tools to transfer large amounts of data into or out of Alibaba Cloud OSS. You can use OSS buckets of the Standard storage class to store image, audio, and video files for apps and large 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.

Terms

- storage class

The storage class of an object. OSS provides the following storage classes to cover various data storage scenarios from hot data to cold data: Standard, IA, Archive, and Cold Archive. OSS Standard storage provides highly reliable, highly available, and high-performance object storage services that can support frequent data access. OSS IA storage is suitable for storing long-lived, but less frequently accessed data (an average of once or twice per month). OSS Archive storage is suitable for long-term storage (at least six months) of infrequently accessed data. OSS Cold Archive storage is suitable for storing extremely cold data with ultra-long lifecycles. For more information, see [Introduction to storage classes](#).

- bucket

The container used to store objects in OSS. Every object is contained in a bucket. You can configure various attributes of a bucket, including the region, ACL, and storage class. You can create buckets of different storage classes to store data based on your requirement. For more information about how to create a bucket, see [Create buckets](#).

- object

The fundamental entities stored in OSS. Objects are also known as files. An object has metadata, data, and key. The key is the unique object name in a bucket. Object metadata describes the attributes of an object such as the last modification time and the object size. You can also specify custom metadata of an object.

- region

The physical location of an OSS data center. You can select the region where your buckets are created based on the cost and the location where the requests come from. For more information, see [Regions and endpoints](#).

- endpoint

The domain name used to access OSS. OSS provides external services by using HTTP RESTful API operations. Different regions use different endpoints. For the same region, access over the internal network or over the Internet also uses different endpoints. For more information, see [Regions and endpoints](#).

- AccessKey pair

The key pair composed of an AccessKey ID and an AccessKey secret. The AccessKey pair is used to perform access identity verification. OSS uses an AccessKey pair, which includes an AccessKey ID and an AccessKey secret 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 the signature string and for OSS to verify the signature string. The AccessKey secret must be kept confidential. For more information about how to obtain an AccessKey pair, see [Create an AccessKey pair](#).

Related services

After you upload your data to OSS, you can use other Alibaba Cloud products and services to perform operations on the data.

The following services are frequently used with OSS:

- **Image Processing (IMG):** a service that allows you to perform various operations on images stored in OSS, including format converting, resizing, cropping, rotating, and adding watermarks. For more information, see [Image processing](#).
- **Elastic Compute Service (ECS):** a cloud computing service that offers elastic and efficient computing capability. For more information, see [ECS product homepage](#).
- **Alibaba Cloud CDN:** a content delivery service that caches the resources on the original site to edge nodes in different regions for quick access. For more information, see [CDN product homepage](#).
- **E-MapReduce:** a big data processing system solution built on ECS. It is based on Apache Hadoop and Apache Spark to facilitate data analysis and processing. For more information, see [E-MapReduce product homepage](#).
- **ApsaraVideo for Media Processing:** a service that converts an audio or video file into another or multiple audio or video files. ApsaraVideo for Media Processing helps you to produce audios and videos that are suitable for PCs, TVs, and mobile devices. On the basis of deep learning of large amounts of data, ApsaraVideo for Media Processing performs multi-model analysis based on the content, text, audio, scenario of audio or video files. It is able to intelligently detect, understand, and edit content. For more information, see [ApsaraVideo for Media Processing product homepage](#).

Manage OSS

- Manage OSS in 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 [Console User Guide](#).

- Manage OSS through APIs or SDKs

OSS provides RESTful API operations and SDKs in multiple programming languages to facilitate secondary development. For more information, see [API Reference](#) and [SDK Reference](#).

- Manage OSS by using tools

OSS provides various tools for you to manage your OSS resources. For more information, see [Tools](#).

OSS pricing

Traditional storage providers require you to purchase a predetermined amount of storage and network transfer capacity. If the amount and capacity is exceeded, your service is disabled or you are charged excess fees. If you do not use the full capacity, you still pay as though you have used it all.

OSS charges you only for what you actually use, without excess fees. As your business grows, you can enjoy the cost advantages of the flexible infrastructure from Alibaba Cloud, which adapts to meet your ever-changing requirements.

For more information about OSS pricing, see [OSS pricing page](#). For more information about OSS billing methods, see [Billing items and methods](#).

Learning path

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

2. Benefits

OSS provides secure, cost-effective, and highly reliable services for you to store large amounts of data in the cloud. This topic compares OSS with the traditional user-created server storage to show the benefits of OSS.

Advantages of OSS over user-created server storage

Item	OSS	User-created server storage
Reliability	<p>OSS is the core infrastructure of data storage for Alibaba Group. This service is highly available and reliable. OSS has provided support during peak hours of Double 11. OSS features a multi-redundant architecture, which provides reliable data storage. In addition, OSS is designed based on a high-availability architecture, which eliminates single point of failures (SPOFs) and ensures the continuity of data-based services.</p> <ul style="list-style-type: none">• Provides 99.995% availability or service continuity (designed for).• Provides data durability of at least 99.9999999999% (twelve 9s).• Automatically expands capacities without affecting your services.• Automatically stores multiple copies of data for backup.	<ul style="list-style-type: none">• Prone to errors due to low hardware reliability. If a disk has a bad sector, data may be lost.• Manual data restoration is complex and requires a lot of time and technical resources.

Item	OSS	User-created server storage
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, IP address blacklist or whitelist, fine-grained permission control, log audit, and Write Once Read Many (WORM) policies. • Provides resource isolation mechanisms for multiple tenants and supports zone-disaster recovery. • Obtains a number of compliance certifications, including certifications from institutions such as SEC and FINRA. This way, the data security and compliance requirements of enterprises can be met. 	<ul style="list-style-type: none"> • Additional scrubbing devices and black hole policy-related services are required. • A separate security mechanism is required.
Cost	<ul style="list-style-type: none"> • Provides multi-line BGP backbone networks with sufficient bandwidth resources. Upstream traffic is free of charge. • Requires no maintenance staff or hosting fees. 	<ul style="list-style-type: none"> • Storage space is limited by hard disk capacity. The storage space must be manually resized. • Single-line or double-line access is slow, and bandwidth is limited. The bandwidth must be manually resized during peak hours. • Professional maintenance staff is required, which results in high costs.
Intelligent storage	<p>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 seamlessly integrates with the Hadoop ecosystem and Alibaba Cloud services such as Function Compute, E-MapReduce, Data Lake Analytics, Batch Compute, MaxCompute, and Database Backup to manage data of enterprises.</p>	<p>Data processing capabilities must be purchased and separately deployed.</p>

More 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 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.
 - Streaming writes and reads are supported, which is suitable for business scenarios where you must simultaneously read and write videos and other large objects.
 - Lifecycle management is supported. You can configure lifecycle rules to delete multiple expired objects or convert the storage class of expired objects to cost-effective Infrequent Access (IA), Archive, or Cold Archive.
- Powerful and flexible security mechanisms
 - Flexible authentication and authorization mechanisms are available. OSS provides STS and URL authentication and authorization. OSS also supports IP address blacklists or whitelists, hotlink protection, and RAM.
 - OSS provides user-level resource isolation. You can also use the multi-cluster synchronization service.
- Data redundancy mechanism

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

 - Operations on objects in OSS are highly consistent. For example, when you receive an upload or copy success response, the uploaded object can be immediately read, and the copies of the object are written to multiple devices for redundancy.
 - To ensure complete data transmission, OSS checks for errors when packets are transmitted between the client and the server. To do so, OSS calculates the checksum of the network traffic packets.
 - The data redundancy mechanism of OSS can prevent data loss when two storage devices 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 the reliability and availability of data.
 - 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 reconstructs and repairs the corrupted data by using the other copies.
- Rich and powerful value-added services
 - IMG: supports format conversion, thumbnails, cropping, watermarking, 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.
 - Accelerated Internet access: provides transfer acceleration that accelerates uploads and downloads across provinces and continents and improves user experience. For more information, see [Transfer acceleration](#).

- Accelerated content delivery: uses OSS as the origin with CDN to improve user experience when the same object is repeatedly downloaded.

3.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. OSS supports various devices. Websites and mobile applications can directly read or write OSS data. OSS supports file writing and streaming writing.

□

Dynamic and static content separation for websites and mobile applications

By using high bandwidth, you can concurrently download a large amount of data from OSS over the Internet. OSS provides [Transfer acceleration](#) that accelerates uploads and downloads across provinces and continents and improves user experience. OSS can also be combined with CDN to deliver the static content to edge nodes. OSS uses data cached on CDN edge nodes to accelerate the simultaneous download of an object in the same region.

□

Cloud data processing

After an object is uploaded to OSS, you can use ApsaraVideo for Media Processing and IMG to process the object on the cloud.

□

4.Features

This topic lists the common features of OSS.

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 lists the features provided by OSS.

Scenario	Description	Reference
Upload objects	Before you upload objects to OSS, you must create a bucket in an Alibaba Cloud region to store your objects. After you create a bucket, you can upload objects to the bucket.	<ul style="list-style-type: none"> • Create buckets • Create folders • Simple upload
Search for objects	You can search for objects or folders in a bucket.	Search for objects
View, share, and download objects	You can use the object URLs to view, share, and download objects.	Download objects
Delete objects or folders	You can delete objects or folders in OSS buckets. You can also delete parts generated in multipart uploads to save storage space.	<ul style="list-style-type: none"> • Delete objects • Delete folders • Manage parts
Automatically delete multiple objects at a specified point in time	You can specify and manage the lifecycle of specific objects or all objects in a bucket. For example, you can delete a specific object or convert it to a cheaper storage class after a period of time.	Lifecycle rules
Accelerate upload and download of data	You can use transfer acceleration to improve customer experience and reduces the amount of time required for businesses to deliver services to their customers.	Transfer acceleration
Recover accidentally deleted data	You can use versioning to recover data that is accidentally overwritten or deleted.	Configure versioning
Zone-disaster recovery	OSS uses the multi-zone mechanism to distribute user data across three zones within the same region. Even if one zone becomes unavailable, the data is still accessible. The zone-redundant storage (ZRS) feature is designed to provide 99.999999999 (twelve 9s) durability and 99.995% availability of data over a given year.	Zone-redundant storage

Scenario	Description	Reference
Geo-location disaster recovery	You can use cross-region replication (CRR) to synchronize operations performed on data such as create, update, and delete operations from the source bucket to the destination bucket. This way, you can implement geo-disaster recovery.	Cross-region replication
Data retention compliance	You can create a retention policy to set the data retention period. Your data cannot be deleted by any users within this period.	Retention policy
Control access to data	<p>You can control access to OSS resources by using the following methods:</p> <ul style="list-style-type: none"> • ACL: You can use access control lists (ACLs) to set the access permission of buckets and objects, including public read/write, public read, and private. • Bucket policy: You can use bucket policies to authorize other users to access your OSS resources by using the console. For example, you can authorize RAM users of other Alibaba Cloud accounts to access your OSS resources, and authorize anonymous users to access your OSS resources from specific IP addresses. • RAM policy: You can create RAM policies to control access to buckets and folders. 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 RAM user a temporary access credential that has 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"> • Set the ACL for a bucket • Configure ACL for objects • Bucket Policy • RAM Policy • STS临时授权访问OSS • Configure hotlink protection
Encrypt data	You can use client-side and server-side encryption to encrypt your data before you store the data in OSS.	<ul style="list-style-type: none"> • Client-side encryption • Server-side encryption

Scenario	Description	Reference
Manage data based on the data category	<p>You can use tagging configuration to manage data based on the data category:</p> <ul style="list-style-type: none"> • Configure bucket tagging: You can manage data based on bucket tags. For example, you can list buckets that have specified tags and configure the ACL for buckets that have specified tags. • Configure object tagging: You can manage data based on object tags. For example, you can configure lifecycle rules for objects that have specified tags and configure the ACL for objects that have specified tags. 	<ul style="list-style-type: none"> • Configure bucket tagging • Configure object tagging
Record the access information for OSS resources	You can enable logging to record the detailed access information for OSS resources.	Logging
Access OSS resources by using custom domain names	You can bind a custom domain name to an OSS bucket and then use the custom domain name to access data in the bucket. If you want to use your custom domain name to access OSS by using HTTPS, you can host your SSL certificate in OSS.	<ul style="list-style-type: none"> • Bind custom domain names • Host SSL certificates
Static website hosting	You can configure static website hosting for your bucket and access this static website by using the bucket endpoint.	Static website hosting
CORS	OSS provides Cross-origin resource sharing (CORS) in HTML5. CORS allows client web applications that are loaded in one domain to interact with resources in another domain.	Configure CORS
Obtain content from the origin	You can configure back-to-origin rules to specify whether to obtain source data by using mirroring or redirection. You can configure back-to-origin rules for online data migration and specific request redirection.	Manage back-to-origin configurations
View object metadata	You can use bucket inventories to export the metadata of a specified object including the object size and encryption status.	Bucket inventory
Modify HTTP headers	OSS allows you to configure HTTP headers to customize HTTP request policies such as the cache policy and forced object download policy.	Configure object HTTP headers
View resource usage	You can view real-time information about OSS service usage, such as the status and performance of the basic system operation.	Overview of the monitoring service

Scenario	Description	Reference
Traffic throttling	You can configure single-connection bandwidth throttling for upload, download, and copy operations on OSS to ensure sufficient bandwidth for other applications.	Single-connection bandwidth throttling
Analyze and process data	<p>You can analyze and process data stored in OSS:</p> <ul style="list-style-type: none"> • Image processing (IMG): You can perform operations such as format conversion, cropping, scaling, rotating, watermarking, and style encapsulation on images stored in OSS. • Capture video snapshots: You can capture images from video objects in the H.264 format. • Intelligent Media Management (IMM): OSS uses IMM to support various data analysis and processing operations such as document preview, document format conversion, facial recognition, image analysis, and QR code recognition. 	<ul style="list-style-type: none"> • IMG • Capture video snapshots • IMM
Use tools to manage OSS resources	OSS provides graphical, CLI, file mounting, and FTP tools for you to manage OSS resources.	OSS tools
Use SDKs to manage OSS resources	OSS provides SDKs that use a variety of programming languages to facilitate secondary development.	SDK sample code

5. Terms

This topic describes several basic terms used in OSS.

bucket

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

- OSS does not have the hierarchical structure of directories and subfolders as in a file system. All objects are directly related to their corresponding buckets.
- You can create multiple buckets.
- A bucket name must be globally unique within OSS. Bucket names cannot be changed after the buckets are created.
- A bucket can contain an unlimited number of objects.

The naming conventions for buckets:

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


object

The fundamental entities stored 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 folders 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 uniquely 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 modification 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. During the lifecycle, content can be appended only to objects created by using append upload. If you want to modify an object, you must upload a new object that has the same name as the existing object to replace the existing object.

The naming conventions for objects:

- The name can contain only UTF-8 characters.
- The name must be 1 to 1,023 bytes 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

ObjectKey, Key, and ObjectName are the same in different SDKs. They all indicate the object names that you must specify when you perform related operations on an 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 an OSS data center. When you create a bucket, you can select a region based on the cost and request source. In most cases, users can access OSS data centers from closer regions faster than centers from farther regions. 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 data center. Regions are configured for buckets instead of objects.

endpoint

The domain name 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 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 that includes an AccessKey ID and an AccessKey secret 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 the signature string and for OSS to 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 RAM to assign the AccessKey pair to a third party.
- The bucket owner uses STS to assign the AccessKey pair to a third party.

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

Strong consistency

A feature requires that object operations in OSS be atomic, which indicates that operations can only either succeed or fail without intermediate states. To ensure that users can access only complete data, OSS does not return corrupted or partial data.

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

Data redundancy mechanism

OSS uses the data redundancy mechanism that is based on erasure coding and multiple replicas to store copies of each object in multiple devices across different facilities within the same region. This way, data reliability and availability are ensured even if hardware failures occur.

- Operations on objects in OSS are highly consistent. For example, when you receive an upload or copy success response, the uploaded object can be read immediately, and the copies of the object are written to multiple devices for redundancy.
- To ensure complete data transmission, OSS checks for errors when packets are transmitted between the client and the server. To do so, it calculates the checksum of the network traffic packets.

- The data redundancy mechanism of OSS can prevent data loss when two storage devices 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 the reliability and availability of data.
 - 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 reconstructs and repairs the corrupted data by using the other copies.

Comparison between OSS and file systems

Item	OSS	File system
data model	OSS is a distributed object storage service that uses a key-value pair format.	A file system uses a typical tree index structure.
Data retrieval	Objects are retrieved based on unique object names (keys). For example, 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 a.jpg. Therefore, similar amounts of resources are consumed regardless of which object you access.	To access a file named test1/test.jpg, you must first access the test1 directory and then search for the test.jpg file in this directory.
Benefits	OSS supports sporadic bursts of access from multiple 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 folder operations such as renaming, deleting, and moving folders.

Item	OSS	File system
Disadvantages	<p>Objects stored in OSS cannot be modified. A specific operation must be called to append an object, and the generated object is different from objects uploaded by using other methods. To modify even a single byte, you must upload the entire object again.</p> <p>You can simulate similar functions of a file system in OSS, but such operations are costly. For example, if you want to rename the test1 directory as test2, OSS must copy all objects whose names start with test1/ to generate 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. Creating more files and directories in the file system consumes greater resources and leads to longer processing time.</p>

We recommend that you do not map operations on OSS objects to file systems because it is inefficient. If you attach OSS as a file system, we recommend that you only add new files, delete files, and read files. You can make full use of OSS advantages, such as the capability to process and store large amounts of unstructured data such as images, videos, and documents.

The following table lists the differences of some terms between OSS and file systems.


OSS	File system
Object	file
Bucket	home directory
Region	None
Endpoint	None
AccessKey	None
None	multilevel directory
GetService	obtain the list of home directories
GetBucket	obtain the list of files
PutObject	add a file
AppendObject	append data to an existing file

OSS	File system
GetObject	read a file
DeleteObject	delete a file
None	modify file content
CopyObject (same destination and source buckets)	modify file attributes
CopyObject	copy a file
None	rename a file

6.Limits

This topic describes the performance metrics and limits of OSS.

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

Item	Limit
Bandwidth	<p>Default bandwidth limit: 10 Gbit/s in mainland China regions and 5 Gbit/s in regions outside mainland China. When this limit is reached, requests are throttled.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Note When requests are throttled, the response header contains <code>x-oss-qos-delay-time: number</code>. The returned <code>number</code> indicates the time period during which requests are throttled. Unit: ms.</p> <ul style="list-style-type: none"> • For an upload request, the exact time period during which requests are throttled is returned. • For a download request, the estimated time period during which requests are throttled is returned. The time is estimated based on the throttling status and the object size. </div> <p>If you require a greater bandwidth (10 Gbit/s to 100 Gbit/s) or queries per second (QPS), for example, if you have offline big data processing requirements, contact the technical support.</p>
QPS	<p>The limit of the total QPS for a single account is 10,000. The actual values that can be achieved are different in the different read and write modes:</p> <ul style="list-style-type: none"> • Sequential read/write: 2,000 <p>If you upload a large number of objects with sequential prefixes such as timestamps and letters in the object names, multiple object indexes may be stored in a single partition. If excessive requests are sent to query these objects, the responsiveness may be slow. We recommend that you do not upload a large number of objects by using sequential prefixes. 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/write: 10,000 <p>If you require a greater QPS, contact the 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. • After a bucket is created, its name, region, and storage class cannot be modified. • OSS does not impose limits on the capacity of a single bucket.

Item	Limit
Archive	It takes about one minute to restore data from the frozen status to the readable status.
Upload objects	<ul style="list-style-type: none"> Objects larger than 5 GB cannot be uploaded by using the following methods: console upload, simple upload, form upload, and append upload. Objects larger than 48.8 TB cannot be uploaded by using multipart upload. If you upload an object that has the same name of an existing object in OSS, the new object overwrites the existing object.
Object management	When you change the storage class of objects by calling the CopyObject operation in the console, the object size cannot exceed 1 GB.
Object deletion	<ul style="list-style-type: none"> Deleted objects cannot be recovered. You can delete up to 100 objects at a time from the console. To delete more than 100 objects at a time, you must call API operations or use SDKs.
Domain name binding	<ul style="list-style-type: none"> You must complete real-name verification for your account on the Alibaba Cloud official website. For domain names bound to buckets located within mainland China regions, you must apply for an ICP filing at the Ministry of Industry and Information Technology (MIIT). A single bucket can be bound to up to 100 domain names, but a domain name can be bound only to a single bucket.
Lifecycle	You can configure up to 1,000 lifecycle rules for each bucket.
Back-to-origin rules	<ul style="list-style-type: none"> You can configure up to 20 back-to-origin rules for each bucket. In regions within mainland China and China (Hong Kong), the default QPS of mirroring-based back-to-origin is 2,000, and the default bandwidth is 2 Gbit/s. In regions outside China, the default QPS for mirroring-based back-to-origin is 1,000, and the default bandwidth is 1 Gbit/s.

Item	Limit
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 objects are supported.▪ The object size cannot exceed 20 MB.▪ If you use image rotation, the width or height of the image cannot exceed 4,096 pixels.▪ Each side of the image cannot exceed 30,000 pixels.▪ The source image cannot exceed 250 million pixels in total.◦ Resized images<ul style="list-style-type: none">▪ The product of dimensions cannot exceed 4,096 × 4,096 pixels.▪ Each side of the image cannot exceed 4,096 pixels.• Limits on image styles<p>You can create up to 50 image styles in each bucket. If you require more than 50 styles, contact the technical support.</p>