

Alibaba Cloud

Object Storage Service Security white paper

Document Version: 20201010

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.Optimization of Alibaba Cloud storage services	05
1.1. Overview	05
1.2. Data storage requirement assessment	05
1.3. Alibaba Cloud storage service selection	06
1.4. OSS optimization	07
1.5. Block Storage optimization	08
1.6. Continuous storage optimization	09
2.Security white paper	10
2.1. Overview	10
2.2. Access control	11
2.3. Data encryption	13
2.4. Monitoring and audit	14
2.5. Disaster recovery	15
2.6. Data retention compliance	16
2.7. Other features	17

1. Optimization of Alibaba Cloud storage services

1.1. Overview

This topic describes how to select and optimize Alibaba Cloud storage services to help you meet your data storage requirements and save storage costs.

In general, enterprises and organizations regard data storage as an auxiliary service. Therefore, they do not optimize their storage or clear the storage that is not used after their data is uploaded to the cloud, which results in the huge cost of storage services. According to [a blog post by RightScale](#), about 7% of the cost for cloud services is wasted on storage volumes that are not used and old snapshots that are copies of storage volumes.

Alibaba Cloud provides various flexible data storage solutions for different storage resources, including blocks, files, and objects. These solutions allow you to convert the storage type of your data at any time. This topic describes how to select the Alibaba Cloud storage services that can best meet your data storage requirements with the lowest cost. This topic also describes how to optimize the storage services that you select to achieve a balance among performance, availability, and durability.

1.2. Data storage requirement assessment

Before you optimize your storage, you must understand the performance profile of each service load and measure performance data such as IOPS and throughput.

Alibaba Cloud storage services provide various storage optimization solutions for different scenarios because no solution applies to all scenarios. Therefore, when you assess your storage requirements, select different storage solutions based on service loads.

When you categorize data and determine the storage requirements of each service load, consider the following factors:

- Data size

The total size of data can help you evaluate the storage capacity and costs.

- Data access frequency and response time requirements

Alibaba Cloud provides various storage solutions that are different in prices based on data access frequency and the requirements on response time.

- Requirements on IOPS and throughput

Alibaba Cloud provides different storage types based on the requirements on IOPS and throughput. You can select appropriate storage types based on your requirements on IOPS and throughput to avoid unnecessary costs.

- Importance of data


Critical or regulated data must be stored securely for extended periods.

- Data sensitivity

Highly sensitive data must be protected not only from missing and damages but also from accidental or malicious modification. Data durability and security are equally important as storage costs.

1.3. Alibaba Cloud storage service selection

Select appropriate Alibaba cloud storage services that best meet your requirements on data availability, data durability, and performance.

 **Note** Data availability indicates the ability of a storage service to provide data based on requests. Data durability indicates the annual average expected data loss of a storage service. Performance indicates the IOPS or throughput that a storage service can provide.

Alibaba Cloud provides the following three storage services to meet different requirements: Object Storage Service (OSS), Block Storage, and Apsara File Storage NAS. You can select the storage solution that best meets your requirements.

OSS

OSS is a secure, cost-effective, and highly reliable cloud storage service provided by Alibaba Cloud. It is suitable to store unstructured data such as audio and video. OSS provides the highest level of data durability and availability among Alibaba Cloud storage services. OSS provides the following three storage classes: Standard, Infrequent Access (IA), and Archive. The three storage classes apply to hot data, warm data, and cold data respectively. The colder the data is, the lower the storage cost is, and the higher the cost for accessing the data is. You can easily convert the storage class of your data to optimize your storage costs.

Block Storage

Block Storage is a high-performance, low-latency block storage service for Alibaba Cloud ECS. You can think of a Block Storage device as a physical disk. You can format a Block Storage device and create a file system on it.

Alibaba Cloud provides a variety of **Block Storage devices** for ECS instances, such as cloud disks based on a distributed storage architecture, and local disks located on the physical machines where the ECS instances are hosted. Cloud disks and local disks are described as follows:

- Cloud disks are block-level storage devices provided by Alibaba Cloud for ECS instances. Cloud disks use a triplicate distributed mechanism and feature low latency, high performance, high durability, and high reliability. Cloud disks can be created, resized, and released at any time.
- Local disks are physical disks attached to physical machines that host ECS instances. Local disks provide local storage access capabilities for ECS instances. Local disks are suitable for scenarios where high storage I/O performance and high cost performance for massive storage are required. Local disks feature low latency, high random IOPS and throughput, and high cost performance.

Cloud disks are billed based on their storage capacity. You can use cloud disks as system disks or data disks. Local disks are billed based on their storage capacity. You can use local disks only as data disks. Local disks cannot be purchased separately. Local disks created together with an ECS instance have the same billing method as the ECS instance. For more information about the product types and prices, see the [Block Storage pricing page](#).

Apsara File Storage NAS

Apsara File Storage NAS is a cloud service that provides file storage for compute nodes, including ECS instances, E-HPC nodes, and Alibaba Cloud Container Service for Kubernetes (ACK) nodes. Apsara File Storage NAS is a distributed file system that supports both the NFS and SMB protocols and features shared access, elastic scalability, high reliability, and high performance.

Apsara File Storage NAS provides the four storage types: Extreme NAS, NAS Performance, NAS Capacity, and Infrequent Access.

Summary

OSS and Apsara File Storage NAS allocate storage resources based on the storage that you use. You only pay for the storage resources that you use. However, when you use Block Storage, you are charged for the pre-allocated storage resources regardless of whether you use the resources. Therefore, to maintain a low storage cost while meeting your requirements, it is important to use OSS to the maximum extent and use Block Storage with pre-configured I/O only when it is required for your application.

1.4. OSS optimization

OSS provides various storage management features to help you optimize your storage performance and cost.

You can analyze the access mode of your data and configure [lifecycle rules](#) to automatically convert data that is infrequently accessed to lower-cost storage classes. To manage data stored in OSS more efficiently, you can add tags to objects to classify them and use tags as filtering conditions in lifecycle rules.

[Bucket inventory](#) helps you understand the status of objects in your buckets and simplify and speed up workflows and big data tasks. The bucket inventory feature scans objects in your bucket on a weekly basis, generates an inventory list in the CSV format, and stores the list as an object in the specified bucket. You can specify object metadata to be exported to the inventory list, such as object size and encryption status.

[OSS monitoring service](#) provides metrics to measure the running status and performance of the system. The monitoring service also provides a custom alert service to help you track requests, analyze usage, collect statistics for business trends, and discover and diagnose system problems in a timely manner.

By using the information provided by the preceding storage management features, you can configure lifecycle rules to convert data that is infrequently accessed to low-cost storage classes to realize significant cost savings. For example, you can save up to 40% of your storage cost by converting the storage class of your data from Standard to IA and save up to 70% of your storage cost by converting the storage class of expired data to Archive.

The following table compares the monthly storage cost (including data retrieval fees) of 1 PB of Standard data and IA data. According to the data in the table, if 10% of your data is accessed every month, the IA storage saves 31% of the cost. If 50% of your data is accessed every month, the IA storage saves 20% of the cost. Even if 100% of your data is accessed every month, the IA storage still saves 6% of the cost.

Total size of data	Percentage of accessed data	Cost of Standard storage (CNY)	Cost of IA storage (CNY)	Saved cost
1 PB	10%	125,829	87,294	31%
1 PB	50%	125,829	100,925	20%
1 PB	100%	125,829	117,965	6%

To further optimize storage and data retrieval costs, OSS provides the **OSS Select** feature. In general, an object in OSS is accessed as a whole regardless of the object size. OSS Select allows you to use simple SQL statements to retrieve objects. Therefore, your application does not need to use computing resources to scan and filter the data in objects. OSS Select can increase query performance by four times and reduce query cost by 80%. OSS also supports the retrieval of IA and Archive objects. Therefore, you can find the data to analyze without performing data retrieval operations. By using OSS Select, you can reduce query cost and obtain more data insights.

1.5. Block Storage optimization

When you use Block Storage, you are charged for the pre-configured storage capacity even if the disk is not attached or only few write operations are performed on the disk. Therefore, to optimize the performance and cost of Block Storage, you must regularly monitor and identify cloud disks that are underused, overused, and not attached, and adjust the capacity of these disks to meet actual requirements.

Delete cloud disks that are not attached or used

The simplest way to reduce storage cost is to find and delete cloud disks that are not attached to ECS instances. If a cloud disk is not released when the ECS instance to which the disk is attached is stopped or terminated, the cloud disk is not automatically deleted and continues to incur fees. In this case, you must manually delete the cloud disk. You can also check whether read and write operations are performed on a cloud disk in the past few weeks. If a cloud disk in non-production environments is not used for several weeks or not attached to a ECS instance for one month, we recommend that you delete the disk in a timely manner.

Adjust cloud disk capacity

For a overused cloud disk, you can scale up the disk online or offline to increase the capacity of the disk. For enhanced SSDs (ESSDs), you can upgrade the performance level (PL) of the disk online to meet your requirements on performance and capacity.

You can downgrade the PL of a pay-as-you-go ESSD online to reduce storage capacity and cost.

You can also re-initialize a cloud disk to restore the disk to the state when it was created.

Delete old snapshots

If you create an automatic snapshot policy that takes snapshots on a daily or weekly basis, a large number of snapshots are created and stored. You must regularly clean up unnecessary snapshots to reduce storage costs. You can set a retention period for snapshots in automatic snapshot policies to automatically delete snapshots that exceed the retention period. The deletion of snapshots does not affect Block Storage.

1.6. Continuous storage optimization

To maintain a storage architecture that is rational in both size and price, you must optimize your storage continuously. You must optimize your storage every month to use your storage cost more efficiently. You can simplify the optimization in the following methods:

- Establish a mechanism to optimize your storage and configure storage policies continuously.
- Use monitoring services and bills to monitor your storage costs.
- Use object tags and lifecycle rules to continuously optimize your storage during the entire lifecycle of your data.

Storage optimization is a process in which you continuously evaluate the change of your data storage requirements and select the most cost-efficient storage solutions. For OSS, you can configure lifecycle rules to automatically convert data that is infrequently accessed to lower-cost storage classes. For Block Storage, you can monitor the storage usage, adjust the capacity of underused or overused cloud disks, and delete expired snapshots and disks that are not attached to ECS instances to avoid costs on storage resources that are not used. To simplify storage optimization, you can set up a monthly plan for storage optimization tasks and use the various storage management functions provided by OSS to monitor storage costs and evaluate resource usage.

2.Security white paper

2.1. Overview

Alibaba Cloud Object Storage Service (OSS) provides rich security capabilities and supports various security features, including server-side encryption, client-side encryption, hotlink protection based on Referer whitelists, fine-grained access control, log audit, and retention policies based on WORM. OSS provides complete security protection for your data stored in Alibaba Cloud to meet your security and compliance requirements on enterprise data.

OSS is the only cloud service in China that has passed the audit and certification of Cohasset Associates and can meet specific requirements for electronic data storage. OSS buckets configured with retention policies can be used for business that is subject to regulations such as SEC Rule 17a-4(f), CFTC Rule 1.31(c)-(d), and FINRA Rule 4511(c). In addition, OSS has passed the following compliance certification:

- ISO9001, ISO20000, ISO22301, ISO27001, ISO27017, ISO27018, ISO29151, and ISO27701
- BS10012
- CSA STAR
- PCI DSS
- C5
- MTCS
- GxP
- TPN
- Trusted cloud service authentication
- SOC 1/2/3 reports

This topic describes the security capabilities provided by OSS, including the following features:

Access control	OSS provides access control list (ACL), RAM and bucket policies, and hotlink protection based on Referer whitelists to control and manage access to your OSS resources.
Data encryption	OSS provides server-side encryption, client-side encryption, and encrypted transmission based on SSL or TLS to protect data from potential security risks on the cloud.
Monitoring and audit	OSS allows you to store and query access logs to meet your requirements on the monitoring and audit of enterprise data.
Disaster recovery	OSS supports zone-redundant storage and cross-region replication to provide disaster recovery capabilities for data centers in a same region or across multiple regions.
Data retention compliance	OSS supports the Write Once Read Many (WORM) strategy that prevents an object from being deleted or overwritten for a specified period of time. This strategy is applicable to business under the regulations of the U.S. Securities and Exchange Commission (SEC) and Financial Industry Regulatory Authority, Inc. (FINRA).

Other features

OSS provides versioning to prevent data from being accidentally deleted or overwritten. If your bucket is attacked or used to share illegal content, OSS moves the bucket to a sandbox to prevent your other buckets from being affected.

2.2. Access control

OSS provides access control list (ACL), RAM and bucket policies, and hotlink protection based on Referrer whitelists to control and manage access to your OSS resources.

Read and write permissions

OSS provides access control lists (ACLs) for you to control access permissions. ACLs are policies that grant users access permissions on buckets and objects. You can set the bucket or object ACL when creating a bucket or uploading an object. You can also modify the ACL of a created bucket or an uploaded object at any time.

- **Bucket ACL**

Bucket ACLs are used to control access to buckets. The following table describes the ACLs that you can configure for a bucket.


ACL	Description	Access control
public-read-write	Public read/write	All users, including anonymous users, can read and write objects in the bucket. Fees incurred by such operations are paid by the owner of the bucket. Exercise caution when you configure this ACL.
public-read	Public read	Only the bucket owner can perform write operations on objects in the bucket. Other users, including anonymous users, can perform only read operations on objects in the bucket.
private	Private	Only the bucket owner or authorized users can read and write objects in the bucket. Other users, including anonymous users cannot access the objects in the bucket without authorization.

- **Object ACL**

Object ACLs are used to control access to objects. The following table describes the ACLs that you can configure for an object.

ACL	Description	Access control
public-read-write	Public read/write	All users, including anonymous users, can read and write the object.

ACL	Description	Access control
public-read	Public read	Only the object owner or authorized users can read and write the object. Other users, including anonymous users, can only read the object.
private	Private	Only the object owner or authorized users can read and write the object. Other users, including anonymous users, cannot access the object.
default	Inherited from the bucket	The ACL of the object is the same as that of the bucket that stores the object.

 **Note** By default, the ACL of an object is inherited from the bucket. The ACL of an object takes precedence over the ACL of the bucket that stores the object. Example: If the ACL for an object is set to public-read, all authenticated and anonymous users can read the object regardless of the bucket ACL.

For more information, see [ACL](#).

RAM policies based on users

Resource Access Management (RAM) is a resource access control service provided by Alibaba Cloud. You can configure RAM policies based on the responsibilities of users. You can manage users by configuring RAM policies. For users such as employees, systems, or applications, you can control which resources are accessible. For example, you can create a RAM policy to grant users read permissions on only some objects in a bucket.

A RAM policy is in the JSON format. You can describe a RAM policy by specifying the Action, Effect, Resource, and Condition fields in the Statement field. You can configure multiple Statement fields in a RAM policy to implement flexible authorization. For more information, see [Implement access control based on RAM policies](#).

Bucket policies based on resources

Bucket policies provide resource-based authorization for users. Compared with RAM policies, bucket policies can be configured in the OSS console. In addition, the bucket owner can grant other users permissions to access OSS resources.

By configuring bucket policies, you can authorize RAM users under other Alibaba Cloud accounts to access your OSS resources or authorize anonymous users to access your OSS resources from specific IP addresses. For more information, see [Use bucket policies to authorize other users to access OSS resources](#).

Hotlink protection based on Referrer whitelists

OSS is a pay-as-you-go service. To prevent additional fees caused by unauthorized access to the data in your bucket, you can configure hotlink protection for your buckets based on the Referrer field in HTTP and HTTPS requests.

You can configure a Referer whitelist to allow only requests from specified domain names or HTTP and HTTPS requests that contain the Referer header to access your OSS resources. Hotlink protection can prevent the data in public read or public read/write buckets from hotlinking to protect your legal rights. For more information, see [Configure hotlink protection](#).

2.3. Data encryption

OSS provides server-side encryption, client-side encryption, and encrypted transmission based on SSL or TLS to protect data from potential security risks on the cloud.

Server-side encryption

OSS supports server-side encryption for uploaded data. When you upload data, OSS encrypts the data and stores the encrypted data. When you download data, OSS decrypts the data and returns the original data. A header is added to the response to declare that the data is encrypted on the server.

OSS uses server-side encryption to protect static data. This method is suitable for scenarios where high security or strong compliance is required for object storage. Examples include the storage of deep learning samples and online collaborative documents. You can choose either of the following methods to implement server-side encryption depending on how you choose to manage the encryption keys:

- Server-side encryption that uses CMKs stored in KMS (SSE-KMS)

When you upload an object, you can use a specified CMK ID or the default CMK stored in KMS to encrypt and decrypt large amounts of data. This method is cost-effective because you do not need to send user data to the KMS server through networks for encryption and decryption.

KMS is a secure and easy-to-use management service provided by Alibaba Cloud. KMS ensures the privacy, integrity, and availability of your keys at minimal cost and allows you to securely and conveniently use keys. You can develop encryption and decryption solutions that best suit your needs. You can view and manage your keys in the KMS console.

KMS encrypts data based on AES-256 and stores and manages CMKs used to encrypt data keys. KMS also generates data keys that can be used to encrypt and decrypt large amounts of data. Envelope encryption provided by KMS can protect your data and corresponding data keys from unauthorized access. You can use the default CMK stored in KMS or generate a CMK by using your BYOK materials or BYOK materials provided by Alibaba Cloud.

- Server-side encryption that uses OSS-managed keys (SSE-OSS)

This encryption method is an attribute of objects. OSS server-side encryption uses AES-256 to encrypt objects with different data keys. CMKs used to encrypt data keys are rotated regularly. This method is suitable to encrypt and decrypt multiple objects at a time.

In this method, data keys are generated and managed by OSS. To perform server-side encryption on an object, you can set the default server-side encryption method of the bucket to KMS without specifying a CMK ID. When sending a request to upload an object or modify the metadata of an object, you can include the `x-oss-server-side-encryption` field in the request and set its value to `AES256`.

For more information, see [Server-side encryption](#).

Client-side encryption

Client-side encryption is performed to encrypt objects on the local client before they are uploaded to OSS. When you use client-side encryption, you must ensure the integrity and validity of the CMK. When you copy or migrate encrypted data, you must ensure the integrity and validity of the object metadata related to client-side encryption.

In client-side encryption, a random data key is generated for each object to perform symmetric encryption on the object. The client uses a CMK to encrypt the random data key. The encrypted data key is uploaded as a part of the object metadata and stored in the OSS server. When an encrypted object is downloaded, the client uses the CMK to decrypt the random data key and then uses the data key to decrypt the object. The CMK is used only on the client and is not transmitted over the network or stored in the server, which ensures data security.

You can use CMKs managed in one of the following ways:

- Use KMS-managed CMKs

If you use KMS-managed CMKs for client-side encryption, you need only to specify the CMK ID when uploading objects instead of providing the client with a data key.

- Use customer-managed CMKs

To use this method for client-side encryption, you must generate and manage CMKs by yourself. When you implement client-side encryption on an object to upload, you must upload a symmetric or asymmetric CMK to the client.

For more information, see [Client-side encryption](#).

Encrypted transmission based on SSL or TLS

OSS supports access through HTTP and HTTPS. You can configure a bucket policy to allow only access through HTTPS (TLS) for better security in data transmission. Transport Layer Security (TLS) is a cryptographic protocol that provides end-to-end communications security over networks. For more information, see [Use bucket policies to authorize other users to access OSS resources](#).

2.4. Monitoring and audit

OSS allows you to store and query access logs to meet your requirements on the monitoring and audit of enterprise data.

Logging

When you access OSS, a large number of access logs are generated. After you enable and configure logging for a bucket, OSS generates an object based on the predefined naming conventions. Access logs are generated on an hourly basis and written to the specified bucket as objects. You can configure lifecycle rules for the specified destination bucket to convert the storage class of the log objects to Archive. This way, these log objects can be retained for a long time. For more information, see [Logging](#).

Real-time log query

OSS uses Log Service to support real-time log query. In the OSS console, you can query and collect statistics for access logs and audit access in OSS, track exception events, and troubleshoot problems. Real-time log query helps you improve work efficiency and make informed decisions. For more information, see [Real-time log query](#).

Monitoring service

The monitoring service of OSS provides metrics to measure the running status and performance of the system. The monitoring service also provides a custom alert service to help you track requests, analyze usage, collect statistics for business trends, and discover and diagnose system problems in a timely manner. For more information, see [Overview](#).

SDDP

Data stored in OSS may include sensitive information such as personal data, passwords, keys, and sensitive images. You can combine OSS with Sensitive Data Discovery and Protection (SDDP) to better identify, classify, and protect sensitive data. After you authorize SDDP to scan your OSS buckets, SDDP identifies sensitive data in your OSS buckets, classifies and displays sensitive data by risk level, and tracks the use of sensitive data. In addition, SDDP protects and audits sensitive data based on built-in security rules, so that you can query the security status of your data assets in OSS buckets at any time. For more information, see [Sensitive data protection](#).

2.5. Disaster recovery

OSS supports zone-redundant storage and cross-region replication to provide disaster recovery capabilities for data centers in a same region or across multiple regions.

Zone-redundant storage (ZRS)

ZRS distributes user data across three zones within the same region. Even if one zone becomes unavailable, the data is still accessible. The ZRS feature can provide data durability (designed for) of 99.999999999% (twelve 9's) and service availability of 99.995%.

ZRS offers data center-level disaster recovery capabilities. When a data center is unavailable due to network disconnection, power outage, or other disaster events, OSS can provide highly consistent services. This way, the service is not interrupted and data is not lost during the failover. This meets the strict requirements of key business systems that the recovery time objective (RTO) and the recovery point objective (RPO) must be zero.

ZRS supports the Standard and Infrequent Access (IA) storage classes. The following table compares the two storage classes from different dimensions.

Index	Standard	IA
Data durability (designed for)	99.999999999% (twelve 9's)	99.999999999% (twelve 9's)
Service availability	99.995%	None
Service availability (designed for)	None	99.995%
Minimum billable size of objects	Actual size of objects	64 KB
Minimum storage period	N/A	30 days
Data retrieval fees	None	Based on the size of retrieved data. Unit: GB.
Data access	Real-time access with low latency (within milliseconds)	Real-time access with low latency (within milliseconds)

Index	Standard	IA
Image Processing (IMG)	Supported	Supported

For more information, see [Zone-redundant storage](#).

Cross-region replication

Cross-region replication (CRR) enables the automatic and asynchronous (near real-time) replication of objects across buckets in different OSS regions. Operations such as the creation, overwriting, and deletion of objects can be synchronized from a source bucket to a destination bucket.

CRR can meet the following business requirements:

- **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.

CRR can meet your requirements on geo-disaster recovery and data replication. Objects in the destination bucket are exact replicas of those in the source bucket. They have the same object names, versioning information, object content, and object metadata such as the creation time, owner, user metadata, and object ACLs. CRR can replicate objects that are not encrypted and objects that are encrypted by using SSE-KMS or SSE-OSS at the server side.

For more information, see [Cross-region replication](#).

2.6. Data retention compliance

OSS supports the Write Once Read Many (WORM) strategy that prevents an object from being deleted or overwritten for a specified period of time. This strategy is applicable to business under the regulations of the U.S. Securities and Exchange Commission (SEC) and Financial Industry Regulatory Authority, Inc. (FINRA).

OSS is the only cloud service in China that has passed the audit and certification of Cohasset Associates and can meet specific requirements for electronic data storage. OSS buckets configured with retention policies can be used for business that is subject to regulations such as SEC Rule 17a-4(f), CFTC Rule 1.31(c)-(d), and FINRA Rule 4511(c). For more information, see [OSS Cohasset Assessment](#).

OSS provides strong compliant policies. You can configure time-based retention policies for buckets. After a retention policy is locked, you can read objects from or upload objects to buckets. However, the objects or retention policies within the retention period cannot be deleted. You can delete objects only after their retention period ends. The WORM strategy is suitable for industries such as finance, insurance, health care, and securities. OSS provides the WORM strategy to allow you to build a compliant bucket in the cloud.

For more information, see [Retention policy](#).

2.7. Other features

OSS provides versioning to prevent data from being accidentally deleted or overwritten. If your bucket is attacked or used to share illegal content, OSS moves the bucket to the sandbox to prevent your other buckets from being affected.

Versioning

OSS provides versioning to prevent your data from being accidentally deleted. After versioning is enabled for a bucket, data that is overwritten or deleted in the bucket is stored as a previous version. Versioning allows you to recover objects in a bucket to any previous version and protects your data from being accidentally overwritten or deleted.

Versioning applies to all objects in buckets. After you enable versioning for a bucket, all objects in the bucket are subject to versioning. Each version has a unique version ID. You can upload objects to and list, download, delete, and recover objects in a versioned bucket. You can also suspend versioning for a bucket to stop the generation of new object versions. When versioning is suspended, you can still specify a version ID to download, copy, or delete a previous version of an object. Each version incur storage fees. You can configure lifecycle rules to automatically delete expired versions.

For more information, see [Overview](#).

OSS sandbox

If your bucket is attacked or used to share illegal content, OSS moves the bucket to the sandbox. A bucket in the sandbox can respond to requests. However, the service quality is degraded. The users of your application may be aware of the degradation. In this case, you must bear the cost incurred by the attack.

To prevent your bucket from being moved to the sandbox due to attacks, you can use [Anti-DDoS Pro](#) to prevent DDoS attacks and HTTP floods. To prevent your bucket from being moved to the sandbox due to the distribution of illegal content that involves pornography, politics, and terrorism, we recommend that you activate [Content Moderation](#) to periodically scan your bucket to effectively monitor the distribution of illegal content.

For more information, see [OSS sandbox](#).