

Alibaba Cloud

Object Storage Service White Paper

Document Version: 20220107

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.Security and compliance	06
1.1. Overview	06
1.2. Access control	07
1.3. Data encryption	09
1.4. Monitoring and audit	12
1.5. Disaster recovery	12
1.6. Data retention compliance	14
1.7. Other features	14
2.OSS performance optimization	16
2.1. Overview	16
2.2. Recommended design patterns	16
2.3. Guidelines for application performance optimization	18
3.Overview of Alibaba Cloud storage services	20
3.1. Overview	20
3.2. Object Storage Service	21
3.3. Apsara File Storage NAS	24
3.4. Block Storage	27
3.5. Cloud Paralleled File System	29
3.6. Apsara File Storage for HDFS	30
3.7. Tablestore	31
3.8. Cloud Storage Gateway	33
4.Optimization of Alibaba Cloud storage services	37
4.1. Overview	37
4.2. Data storage requirement assessment	37
4.3. Alibaba Cloud storage service selection	38
4.4. OSS optimization	40

4.5. Block Storage optimization	41
4.6. Continuous storage optimization	41

1. Security and compliance

1.1. Overview

Alibaba Cloud Object Storage Service (OSS) has rich security capabilities that support various security protection features, including server-side encryption, client-side encryption, hot link protection based on Referer whitelists, fine-grained access control, log audit, and retention policies based on Write Once Read Many (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 is accredited by Cohasset Associates in audit and meets 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 is certified with the following standards:

- ISO 9001, ISO 20000, ISO 22301, ISO 27001, ISO 27017, ISO 27018, ISO 29151, and ISO 27701
- BS10012
- CSA STAR
- PCI DSS
- C5
- MTCS
- GxP
- TPN
- Trusted cloud service authentication
- SOC 1, SOC 2, and SOC 3

This topic describes the security capability of OSS that contains the following features:

Access control	OSS provides access control lists (ACLs), Resource Access Management (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 monitoring and auditing enterprise data.
Disaster recovery	OSS has disaster recovery capabilities that support zone-redundant storage (ZRS) and cross-region replication (CRR) for data centers in a same region or across multiple regions.
Data retention compliance	OSS supports the WORM data storage technology mechanism that prevents users from accidentally deleting or tampering with your data. This mechanism 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 one of your buckets is attacked or used to distribute illegal content, OSS automatically moves the bucket to the sandbox to prevent your other buckets from being affected.

1.2. Access control

OSS provides access control lists (ACLs), 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.
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.

ACL	Description	Access control
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 the 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 [Overview](#).

Temporary access authorization based on STS

Compared with RAM policies that provide long-term access control, Security Token Service (STS) provides temporary credentials for users to access resources. By calling STS API operations, you can obtain temporary AccessKey pairs and tokens and send them to temporary users to access the corresponding resources. Permissions obtained from STS are strictly restricted and have time limits. Therefore, even if information is unexpectedly leaked, your system will not be severely compromised.

You can use STS to authorize temporary access to OSS. You can use STS to grant a third-party application or your RAM user an access credential with the custom validity period and permissions specified. For more information, see [Use a temporary access credential provided by STS to access OSS](#).

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 [Configure 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 Referer 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](#).

1.3. Data encryption

Object Storage Service (OSS) provides server-side encryption, client-side encryption, and encrypted transmission based on SSL or Transport Layer Security (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 decrypted data. In addition, 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. You can use this method in scenarios in which additional security or compliance is required, such as the storage of deep learning samples and online collaborative documents. You can choose one of the following methods to implement server-side encryption depending on the mechanism you choose to manage the encryption keys:

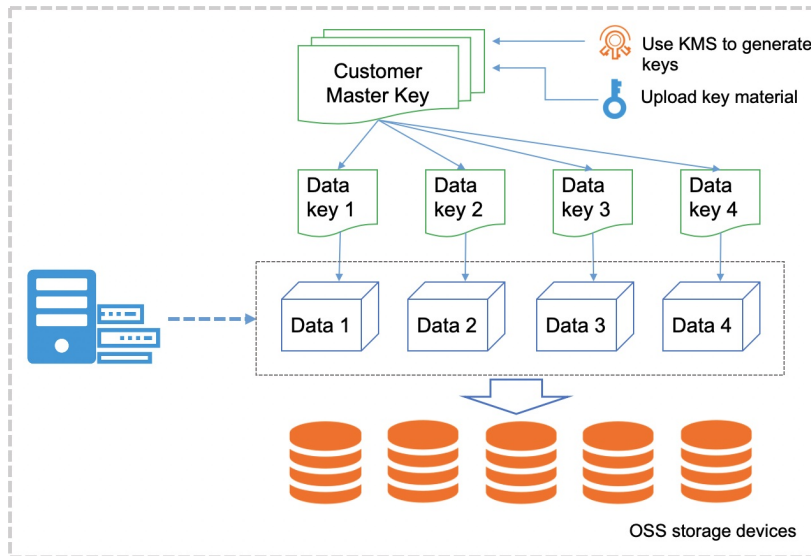
- Server-side encryption that uses customer master keys (CMKs) stored in Key Management Service (KMS) (SSE-KMS)

When you upload an object, you can use a specified CMK ID or the default CMK stored in KMS to encrypt and decrypt the object. This method is applicable to large amounts of data encryption and decryption. This method is cost-effective because you do not need to send data to the KMS server 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. Therefore, you can focus on the development of encryption and decryption functions that best suit your needs. You can view and manage keys in the KMS console.

KMS encrypts data based on AES-256 and stores and manages CMKs that are used to encrypt data keys. KMS also generates data keys that can be used to encrypt and decrypt data. In addition, 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.

The following figure shows the logic of SSE-KMS.



- Multi-level key envelope encryption
- Support for encryption and decryption by using BYOK
- Keys managed by KMS
- A random key generated for each object

- 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 to ensure greater security. This method is suitable for multiple objects encryption and decryption 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 that contains the object to AES-256. You can also include the `x-oss-server-side-encryption` field in the request that is sent to upload an object or modify the metadata of an object and set the field value to `AES-256`.

For more information, see [Server-side encryption](#) in OSS Developer Guide.

Client-side encryption

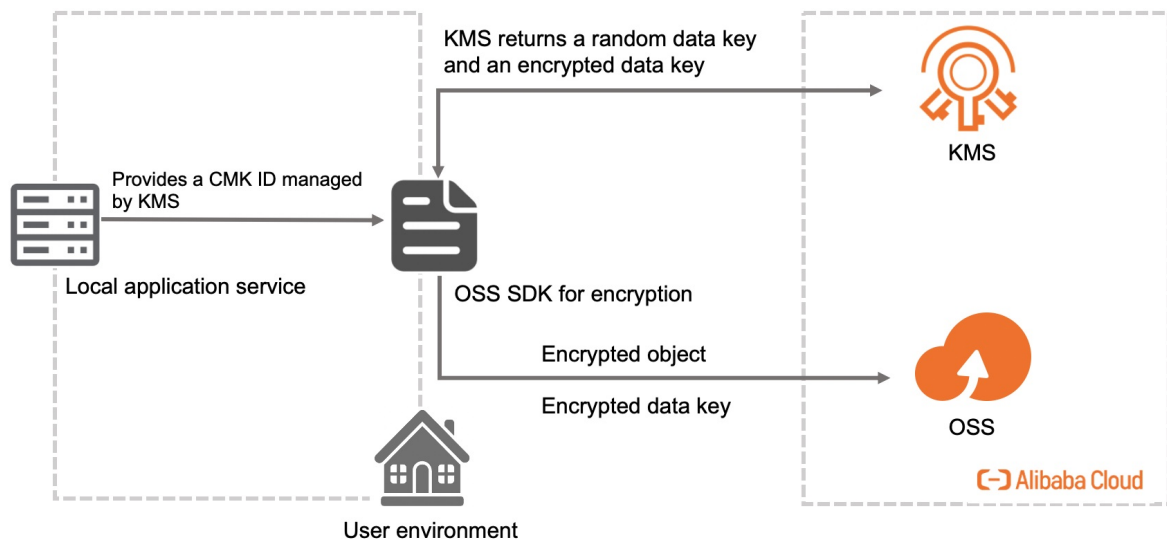
Client-side encryption is performed to encrypt objects on the local client before the objects are uploaded to OSS. When you use client-side encryption, you are responsible for the integrity and validity of the CMKs. When you copy or migrate encrypted data, you are responsible for the integrity and validity 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 generate a data key in random. The encrypted data key is saved 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 manage CMKs in one of the following ways:

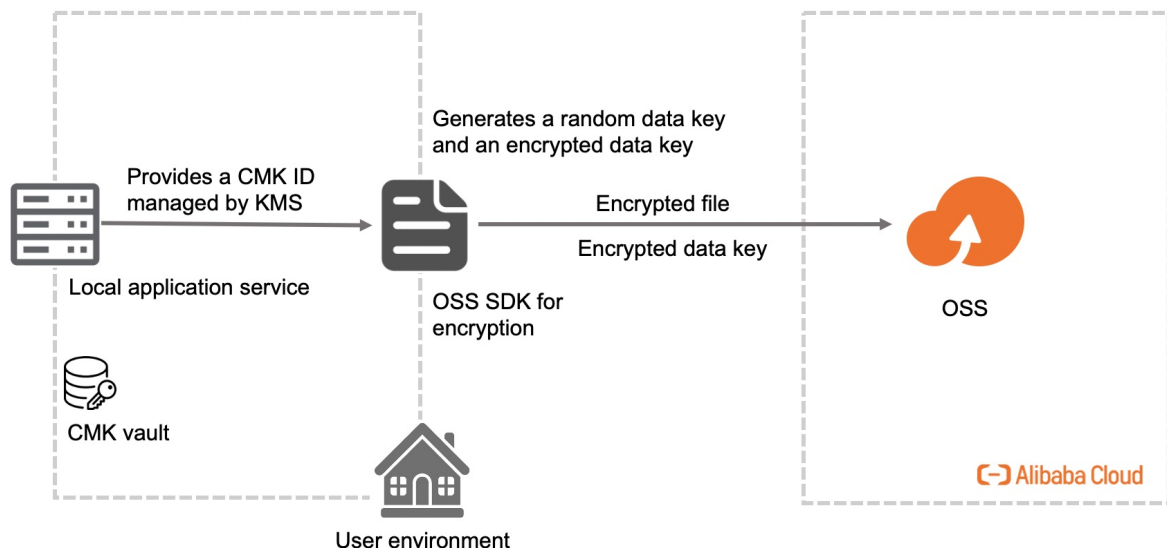
- Use KMS-managed CMKs

If you use KMS-managed CMKs for client-side encryption, you need only to specify the CMK ID when you upload objects without the need to provide the client with a data key. The following figure shows the logic of KMS-managed CMKs for client-side encryption.



- 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. The following figure shows the logic of customer-managed CMKs for client-side encryption.



For more information, see [Client-side encryption](#) in OSS Developer Guide.

Encrypted transmission based on the SSL or TLS protocol

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

1.4. Monitoring and audit

OSS provides logging and real-time log query and supports Inner-ActionTrail for buckets to address your monitoring and auditing needs for enterprise data.

Log storage

When you access OSS, large numbers of access logs are generated. After you enable and configure logging for a bucket, OSS generates log objects every hour in accordance with a predefined naming convention and then stores the access logs as objects in a specified bucket. You can use Alibaba Cloud Log Service or build a Spark cluster to analyze the logs. 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 efficiency and make informed decisions. For more information, see [Real-time log query](#).

Inner-ActionTrail

Alibaba Cloud ActionTrail provides the Inner-ActionTrail feature. This feature allows you to transfer the operation logs of Alibaba Cloud services from ActionTrail to Log Service in near real time for analysis and auditing. ActionTrail records and stores the operations logs of OSS. You can transfer these logs to Log Service for analysis and auditing based on the search and analysis, reporting, alerting, and downstream computing and transferring features. For more information, see [Inner-ActionTrail overview](#).

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 on 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](#).

1.5. 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)
Image Processing (IMG)	Supported	Supported

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

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](#).

1.6. 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](#).

1.7. Other features

Object Storage Service (OSS) provides the versioning feature 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. If your business is prone to distributed denial of service (DDoS) attacks, you can enable the OSS DDoS protection feature for your buckets.

Versioning

OSS provides the versioning feature to prevent your data from being accidentally deleted or overwritten. After versioning is enabled for a bucket, data that is deleted or overwritten in the bucket is stored as a previous version. You can use versioning to recover a previous version of an object that was overwritten or deleted.

Versioning is applied to all objects within a versioned bucket. After you enable versioning for a bucket, versioning applies to all objects in the bucket. OSS generates a unique ID for each version of objects in the bucket. You can upload objects to a versioning-enabled bucket and list, download, delete, and recover the objects. You can also suspend versioning for a bucket to stop OSS from generating new object versions. After versioning is suspended for a bucket, you can still specify a version ID in a request to download, copy, or delete the specified previous version of an object in the bucket. You are charged for each version of objects in your bucket. You can configure lifecycle rules to delete expired versions on a regular basis.

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

OSS sandbox

If your bucket is attacked or used by other users 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 are charged the fees caused by the attack.

To prevent your bucket from being moved to the sandbox due to attacks, we recommend that you 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, we recommend that you activate [Content Moderation](#) to periodically scan your bucket to monitor the distribution of illegal content.

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

OSS DDoS protection

OSS DDoS protection is a proxy-based mitigation service that integrates OSS with DDoS protection. When a bucket for which OSS DDoS protection is enabled suffers DDoS attacks, OSS DDoS protection diverts malicious traffic to an Anti-DDoS Pro or Anti-DDoS Premium instance for scrubbing and then redirects normal traffic to the bucket. This way, your business can power through DDoS attacks and continue to function normally.

OSS DDoS protection can be used to protect your buckets from DDoS attacks that involve up to terabytes of traffic per second and millions of queries per second (QPS), and can switch over between Anti-DDoS Pro and Anti-DDoS Origin within a few seconds. These capabilities can prevent attacks such as SYN flood, ACK flood, Internet Control Message Protocol (ICMP) flood, UDP flood, NTP flood, Simple Service Discovery Protocol (SSDP) flood, DNS flood, and HTTP flood attacks. OSS DDoS protection is suitable for scenarios where your business is prone to attacks, ransom-driven attacks, click farming, and fraudulent traffic.

For more information, see [OSS DDoS protection](#) in OSS Developer Guide.

2.OSS performance optimization

2.1. Overview

When you develop applications that need to upload data to or download data from Object Storage Service (OSS), you can refer to the best practices described in this topic to optimize the performance of your applications.

OSS can handle thousands of requests sent from your application each second when the application upload data to or download data from OSS. OSS sorts objects by their names in alphabetical order and stores objects by partitions. Partitions are split based on the size of objects stored in OSS and the QPS of requests sent to access the objects. OSS allows your application to send at least 3,500 PUT/COPY/POST/DELETE requests and 5,500 GET/HEAD requests per second to access a partition. For example, you can configure 10 prefixes for objects that you want to upload to a bucket to store the objects in 10 different partitions. This way, your application can send 35,000 PUT requests per second to write data to the bucket.

Some data lake applications that use OSS as data sources may need to scan millions or billions of objects and query petabytes of data in OSS. These applications can fully use the bandwidth of each ECS instance, which can be up to 100 Gbit/s, to access OSS. OSS is a distributed storage service that supports parallel access. Therefore, you can aggregate the bandwidth of multiple ECS instances used by an application to provide a throughput of terabits per second to access OSS.

You can use OSS to store small files for latency-sensitive applications, such as social media and messaging applications. This way, the applications can access the small files with a latency of 100 to 200 milliseconds.

You can also use other features provided by OSS to improve the performance of applications that use different architectures. For example, you can use Alibaba Cloud Content Delivery Network (CDN) to cache objects stored in OSS. When a large number of users in the same region download a static object in your bucket at the same time, they can obtain a cached version of the object from the nearest edge node to improve the download speed. If you want to transfer data between the client and OSS over a long distance, you can use the transfer acceleration feature provided by OSS. Transfer acceleration provides an end-to-end acceleration solution by combining intelligent scheduling, protocol stack tuning, optimal route selection, and transmission algorithm optimization with OSS server-side configurations.

2.2. Recommended design patterns

When you design applications that need to upload data to or download data from OSS, we recommend that you use the design patterns described in this topic to optimize the performance of your applications. We also provide guidelines for application performance optimization. You can refer to these guidelines when you design your applications.

Cache frequently-accessed content

In scenarios where a large number of users of your application in the same region need to download a static object from OSS at the same time, you can cache the object on the edge nodes of Alibaba Cloud Content Delivery Network (CDN). This way, the users can obtain a cached version of the object from the nearest edge node to improve download performance. You can use caches to reduce latency and improve data transmission speed. Applications that use caches send fewer requests to OSS, which reduces the request fees.

Alibaba Cloud CDN is a distributed network that is built on top of the bearer network and consists of servers distributed in different regions. CDN caches the resources on the origin to edge servers distributed in different regions. Users can obtain cached resources from edge servers that are nearest to their locations instead of the origin. This way, direct access to the origin is reduced.

Establish timeout retry mechanisms for latency-sensitive applications

OSS limits the query per second (QPS) of management operations, such as GetService (ListBuckets), PutBucket, and GetBucketLifecycle. If your application initiates a large number of requests at the same time, HTTP 503 may be returned to indicate that the requests are too frequent. In this case, we recommend that you retry the requests after a few seconds.

OSS restricts the QPS of requests sent by an Alibaba Cloud account to 10,000. If you require a higher QPS, contact the [technical support](#). In addition, OSS also poses a limit for QPS of requests that are sent to a single partition. If you send a large number of requests at the same time to access objects within the same partition, OSS may return HTTP 503 even if the QPS of the requests does not exceed 10,000. To avoid this problem, you can configure different prefixes for objects. This way, OSS creates more partitions to store the objects and supports higher QPS. In this case, you need only to wait and then send your requests again. For more information about how to configure prefixes for objects, see [OSS performance and scalability best practices](#).

If you send a large number of requests that are different in size, such as requests that are larger than 128 MB in size, we recommend that you measure the throughput and retry the slowest 5% requests. In general, the responses to small requests that are smaller than 512 KB in size are returned within tens of milliseconds. If you need to retry GET or PUT requests, we recommend that you retry the requests 2 seconds after the requests are sent. If a request fails multiple times, we recommend that you quit the program and send the request again. For example, you can retry a request 2 seconds after the request is sent and then retry the request again after 4 seconds.

If the requests sent by your application are the same in size and you want the response time of all requests to be consistent, we recommend that you identify and retry the 1% slowest requests. In this case, the response time of the requests can be reduced when the requests are retried.

Send requests in a distributed and concurrent manner for high throughput

OSS is a large-scale distributed storage system. To fully use the throughput of OSS, we recommend that you concurrently send requests to OSS and distribute the requests to multiple OSS service nodes. This way, workload can be distributed to multiple paths over the network.

To increase throughput during data transmission, we recommend that you create multiple threads or instances and initiate multiple requests in the threads or instances to concurrently upload and download data. For some applications, you can initiate multiple requests in different threads or instances to concurrently access OSS. You can determine how to distribute requests based on the architecture of your application and the structure of objects that you want to access.

Before you adjust the number of requests that can be concurrently sent, you must measure the resources that are required for requests. We recommend that you first measure the bandwidth and other resources that are consumed by a single request. This way, you can identify the major resource that is required for requests and then determine the number of concurrent requests based on the resource. For example, if 10% of the CPU is required to process a request at a time, you can send up to 10 concurrent requests.

Use transfer acceleration to speed up upload and download over long distance

OSS uses data centers distributed around the globe to implement transfer acceleration. When a request is sent to your bucket, it is parsed and routed to the data center where the bucket resides over the most optimal network path and protocol. The transfer acceleration feature provides an optimized end-to-end acceleration solution to access OSS over the Internet. For more information, see [Transfer acceleration](#).

You can visit [The Comparison of OSS Direct Data Transfer and Accelerated Data Transfer in Different Regions](#) to compare the access speeds when you use the accelerate and default endpoints to access OSS in different regions. Transfer acceleration fees are calculated separately based on the outbound traffic over the Internet. For example, if you use the accelerate endpoint of a bucket for which transfer acceleration is enabled to download 1 GB of data from the bucket, you are charged transfer acceleration fees and outbound traffic fees for 1 GB.

2.3. Guidelines for application performance optimization

To optimize the performance of your application when the application uses OSS, we recommend that you follow the guidelines described in this topic.

Measure performance indexes

Evaluate the performance of different ECS instance types based on your requirements on CPU and network throughput. For more information about the types of ECS instances, see [Overview](#). In addition, we recommend that you use analysis tools to measure the DNS query time, latency, and data transmission speed.

Distribute requests to multiple connections

It is common to distribute requests to multiple connections for high performance in the design of applications. When you develop a high-performance application, you can use OSS as a large-scale distributed storage system instead of a single storage node such as a traditional storage server. You can send multiple concurrent requests to OSS to improve application performance. You can distribute requests to multiple connections to fully use the bandwidth provided by OSS. OSS does not impose limits on the number of connections to a bucket.

Use range download

You can specify the Range header in a GetObject request to obtain the specified bytes of an object. Only data within the specified range is transferred. You can initiate concurrent requests to obtain data within different ranges of the same object. Instead of sending a single request to obtain an entire object, you can send concurrent requests to obtain an object by range for higher throughput. In addition, if you obtain only data within a small range of a large object, your application takes less time to retry a request when no response is returned for the request. For example, you can set the range size to 512 KB or 1 MB. Note that the boundaries of the byte ranges in different requests must be aligned. For more information, see [GetObject](#).

Increase allowed number of retries

Due to its large scale, OSS may take a long time to respond to a request when the request is sent for the first time. In this case, you can retry the request. You can use OSS SDKs to configure the timeout period and the allowed number of retries for requests based on your requirements.

Keep OSS buckets and ECS instances in the same region

The name of a bucket is globally unique in OSS. You must specify the region of a bucket when you create the bucket. The name and region of a bucket cannot be changed after the bucket is created. To reduce latency and data transmission costs, we recommend that you access an OSS bucket by using ECS instances in the same region as the bucket. For more information, see [Access to OSS resources from ECS instances by using an internal endpoint of OSS](#).

Use transfer acceleration

The transfer acceleration feature provided by OSS is suitable for the regular transmission of large objects of gigabytes or terabytes in size over a long distance. OSS uses data centers distributed around the globe to implement transfer acceleration. When a request is sent to your bucket, it is parsed and routed to the data center where the bucket resides over the most optimal network path and protocol. This way, objects can be transferred over a long distance in a fast and secure manner. For more information, see [Transfer acceleration](#).

You can visit [The Comparison of OSS Direct Data Transfer and Accelerated Data Transfer in Different Regions](#) to compare the access speeds when you use the accelerate and default endpoints to access OSS in different regions.

Use the latest versions of OSS SDKs

[OSS SDKs](#) provide built-in support for multiple performance optimization guidelines. OSS SDKs provide simple API operations and are updated on a regular basis to follow the latest best practices. For example, OSS SDKs support built-in retry logic for network errors on the client and server to adapt to different network conditions. OSS SDKs also provide advanced data transmission API operations that follow best practices for the upload or download of large objects.

3. Overview of Alibaba Cloud storage services

3.1. Overview

Alibaba Cloud provides low-cost, high-reliability, and highly-availability storage services for a wide range of storage resources such as blocks, files, and objects. These services are applicable to a variety of scenarios, including data backup, archiving, and disaster recovery. This topic describes the scenarios, performance, security, operations, and cost models of various Alibaba Cloud storage services to help you select the storage service that best meets your business requirements.

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

Object Storage Service	Object Storage Service (OSS) is a secure, cost-effective, and high-reliability cloud storage service provided by Alibaba Cloud. It enables you to store a large amount of data in the cloud. OSS provides scalable capacity and processing capability and multiple storage classes to cover a variety of data storage scenarios from hot data storage to cold data storage, which helps you minimize your storage costs.
Block Storage	Elastic Block Storage (EBS) is a high-performance, low-latency block storage service for Alibaba Cloud Elastic Compute Service (ECS). It supports random read and write operations. You can use EBS as a physical disk. You can format an EBS device and create a file system on it.
Apsara File Storage NAS	Apsara File Storage Network Attached Storage (NAS) is a cloud service that provides shared access, elastic scalability, high reliability, and high performance. To be more specific, Apsara File Storage NAS is a distributed file system that provides file storage for compute nodes, including ECS instances, Elastic High Performance Computing (E-HPC) nodes, and Container Service for Kubernetes (ACK) nodes. Apsara File Storage NAS is designed based on POSIX and compatible with native operating systems.
Cloud Parallel File System	Cloud Parallel File System (CPFS) is a type of parallel file system provided by Alibaba Cloud. CPFS stores data across multiple data nodes in a cluster and allows multiple clients to simultaneously access data. Therefore, CPFS provides data storage services with high input/output operations per second (IOPS), high throughput, and low latency for large and high-performance computing clusters.
Apsara File Storage for HDFS	Apsara File Storage for Hadoop file system (HDFS) is a file storage service for computing resources such as Alibaba Cloud ECS instances and Container Service for Kubernetes (ACK). Apsara File Storage for HDFS can meet the requirements of distributed computing business such as Hadoop on the performance, capacity, and reliability of distributed storage.
Tablestore	Tablestore is a data storage service developed by Alibaba Cloud to store a large amount of structured data. You can use Tablestore to efficiently query and analyze data. Tablestore can store petabytes of data and provides tens of millions of transactions per second (TPS) within milliseconds of latency.

Cloud Storage Gateway	Cloud Storage Gateway (CSG) is a gateway service that can be deployed at your data center and in Alibaba Cloud. CSG uses Alibaba Cloud OSS buckets as backend storage devices, and provides on-premises and in-cloud applications with standard file services over the Network File System (NFS) and Server Message Block (SMB) protocols. CSG also provides block storage services over the Internet Small Computer Systems Interface (iSCSI) protocol.
-----------------------	--

3.2. Object Storage Service

Object Storage Service (OSS) is a secure, cost-effective, and high-durability cloud storage service provided by Alibaba Cloud. It enables you to store a large amount of data in the cloud. OSS supports a designed durability of at least 99.999999999% (twelve 9's) and a designed availability of at least 99.995%. OSS supports RESTful APIs independent from the console. You can store and access any type of data anytime, anywhere, and from any application.

OSS provides the following storage classes to cover various data storage scenarios from hot data storage to cold data storage: Standard, Infrequent Access (IA), Archive, and Cold Archive.

Standard	Standard storage provides high-durability, high-availability, and high-performance object storage services that can handle frequent data access. Standard storage is ideal for storing images for social networking and sharing, storing data for audio and video applications, large websites, and big data analytics.
Infrequent Access (IA)	IA storage is suitable for data that is accessed infrequently, such as only once or twice a month. IA storage offers a storage unit price lower than that of Standard storage and is suitable for long-term data backup of mobile apps, smart devices, and enterprises. It also supports real-time data access.
Archive	Archive storage is suitable for long-term (at least six months) storage of data that is infrequently accessed. OSS takes up to one minute to restore an Archive object before it can be read. Archive storage is suitable for data that you want to store for a long period of time such as archival data, medical images, scientific materials, and video footage.
Cold Archive	Cold Archive storage is suitable for extremely cold data that you want to store for an extremely long period of time. Examples: data that must be retained for an extended period of time due to compliance requirements, raw data that is accumulated over an extended period of time in the big data and AI fields, media resources that are retained in the film and television industries, and archived videos from the online education industry.

Scenarios

OSS is applicable to the following scenarios:

- Storage and distribution of audio, video, and static website data

Each OSS object can be accessed through a unique HTTP URL, which can be used to distribute data. In addition, OSS buckets can be used as the origins for Alibaba Cloud Content Delivery Network (CDN). The storage space of OSS is not partitioned. Therefore, OSS is ideal to store the data of user-driven and data-intensive websites, such as image and video sharing websites. Various devices, websites, and mobile applications can directly read data from or write data to OSS. You can write data to OSS by uploading files or using streams.

- Static website hosting

As a cost-efficient, high-availability, and highly scalable storage solution, OSS can be used to store static HTML files, images, videos, and client scripts such as JavaScript scripts.

- Data warehouse for computing and analysis

OSS provides high scalability to simultaneously access data from multiple computing nodes and prevent you from being restricted by the performance of a single node.

- Data backup and archiving

OSS provides a highly available, scalable, secure, and reliable solution for the backup and archiving of critical data. You can configure lifecycle rules to automatically convert the storage class of cold data to IA or Archive to reduce storage costs. You can also configure cross-region replication (CRR) rules to automatically replicate data between buckets in different regions in an asynchronous manner in near real-time.

Performance

To achieve the fastest speed when you access an OSS bucket from your ECS instance, the ECS instance must be located within the same region as the bucket. Because of how OSS is designed, the latency at the server side is negligible compared with the network latency. In addition, OSS can support a large number of web applications because it can be scaled based on the storage capacity, number of requests, and number of users. If you use multiple threads, applications, or clients to simultaneously access OSS, the total throughput is scaled to a value far higher than the throughput that can be provided or consumed by a single server.

OSS provides the multiple upload feature to improve the upload speed of objects larger than 5 GB in size. By using multipart upload, you can split an object into multiple data blocks (parts) and upload the parts separately. After all parts are uploaded, OSS combines these parts into a complete object. You can use this method to implement resumable upload. Multipart upload is suitable for scenarios in which network connectivity is poor. If a part fails to be uploaded, you can reupload only the failed part instead of the complete object.

To accelerate data access, many developers use OSS in conjunction with search engines such as OpenSearch or databases such as Tablestore and RDS for MySQL. OSS is used to store data, while search engines or databases are used to store metadata, such as the object names, object sizes, and keywords. Metadata stored in databases can be indexed and queried. You can use OSS in conjunction with search engines or databases to locate and query objects stored in OSS.

OSS provides [transfer acceleration](#) to accelerate the upload and download of large objects from a distance, which allows you to dynamically update objects and accelerate the download of objects that are not frequently accessed. Transfer acceleration provides an end-to-end acceleration solution by combining smart scheduling, protocol stack tuning, optimal route selection, and transfer algorithm optimization with OSS server-side configurations.

OSS supports Alibaba Cloud Content Delivery Network (CDN) to accelerate the download of static hot objects. Alibaba Cloud CDN uses OSS buckets as origins and distributes content from the origins to edge nodes. Alibaba Cloud CDN uses its precise scheduling system to distribute requests to optimal edge nodes, which allows users to quickly access the required content. This way, Internet traffic congestion can be eased and the response time is shortened.

Data durability and service availability

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 can still remain accessible. This mechanism can provide 99.999999999% (twelve 9's) data durability and 99.995% data availability.

You can also enable CRR for a bucket. After CRR is enabled, data is replicated between buckets in different regions in an asynchronous and nearly real-time manner. The source bucket and destination bucket can both provide 99.999999999% (twelve 9's) data durability and 99.995% data availability.

Scalability and elasticity

OSS provides high scalability and elasticity. You may encounter problems if you store excessive files in the same directory in a common file system. In contrast, the total storage capacity of OSS and the capacity of a single bucket are not limited. You can store an unlimited number of objects in a bucket. OSS stores the copies of your data in different servers in the same region. All copies have the same performance provided by the high-performance infrastructure of Alibaba Cloud.

Security

OSS provides a variety of security features such as 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 is the only cloud service in China that meets the audit and certification requirements of Cohasset Associates and the specific requirements for electronic data storage. OSS buckets configured with retention policies can be used for business 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 [Overview](#).

Operations

OSS provides standard RESTful API operations. You can use these API operations to store objects in buckets whose names are globally unique. Each object has a key that uniquely identifies the object in the bucket. OSS does not use folders. All elements are stored as objects. However, you can create a simulated folder by creating an object whose name ends with a forward slash (/), such as folder1/folder2/file.

You can use tools or OSS SDKs that encapsulate the API operations to develop applications. OSS provides [SDKs](#) for more than 10 programming languages, including Java, Python, PHP, Go, Android, and iOS. [ossutil](#) is a high-performance command-line tool that provides a variety of simple commands, such as ls, cp, cat, and config, for you to manage buckets and objects. ossutil supports concurrent upload and is supported by the following operating systems: Windows, Linux, and macOS. In addition, you can use the graphic tool [ossbrowser](#) to perform basic operations, such as browsing objects and uploading or downloading objects and folders. You can also use the [OSS console](#), which is a graphic web application, to manage your OSS resources.

You can use the event notification feature to immediately learn about the operations performed on your OSS resources. For more information, see [Overview](#).

Cost model

OSS fees include storage fees, traffic fees, API operation calling fees, data processing fees. OSS charges fees based on the actual amount of resources used. The fees incurred within an hour are billed in the following hour. Fees are calculated based on the following formula: Fees = Actual usage × Unit price. You can use the subscription billing method for some billing items to reduce costs. Subscription allows you to use resources only after you purchase resource plans. Resource plans are used to deduct fees incurred by resource usage. For more information, see [Billing items and methods](#).

3.3. 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.

Extreme NAS	Extreme NAS is a high-performance file sharing solution that is built on top of the latest generation of network architecture and all-flash storage. The maximum capacity is 256 TB. The bandwidth ranges from 150 Mbit/s to 1,200 Mbit/s. Extreme NAS can provide a constant latency of about 100 microseconds. Extreme NAS is applicable to latency-sensitive business in which a large amount of small files are handled.
NAS Performance	NAS Performance uses solid-state drives (SSDs) as the storage devices, and provides high throughput, high input/output operations per second (IOPS), and low latency for workloads. NAS Performance is a file sharing solution that is applicable to scenarios where high throughput, high concurrency, business scalability, and low read latency are required. You can use NAS Performance if you need to perform frequent read/write operations and have high requirements for response latency.
NAS Capacity	NAS Capacity uses SATA hard disk drives (SATA HDDs) as storage devices and provides high-performance storage space at low costs. NAS Capacity is a file sharing solution that is applicable to cost-sensitive scenarios where high throughput, high concurrency, and business scalability are required. NAS Capacity is more cost-efficient if you do not need to perform frequent read/write operations and do not have high requirements on response latency.
Infrequent Access	You can configure lifecycle rules to transfer data that is infrequently accessed and needs to be stored for long periods from NAS Performance or NAS Capacity to Infrequent Access to reduce costs. For more information, see Implementation of lifecycle management .

Scenarios

Apsara File Storage NAS is applicable to the following scenarios:

- Container storage

You can use containers to build microservices because containers support fast pre-configuration and flexible resource allocation, and can isolate processes. If some containers must access raw data each time they are started, you must create a shared file system for the containers. This way, the containers can access the file system regardless of the instance on which they run. You can use Apsara File Storage NAS as container storage because it provides persistent shared access to files.

- Content management and web services

Apsara File Storage NAS provides high persistence and high throughput. Therefore, you can use Apsara File Storage NAS in content management systems and web servers to store and provide data for websites, online publishing applications, and archiving applications. Apsara File Storage NAS follows the expected file system semantics, file naming conventions, and permissions that are preferred by web developers. You can integrate NAS with web applications and use NAS in websites, online publishing applications, and archiving applications.

- Enterprise applications

Apsara File Storage NAS provides high scalability, elasticity, availability, and persistence. Therefore, you can use Apsara File Storage NAS as storage solutions for your enterprise applications and applications delivered as services (ADaaS). Apsara File Storage NAS provides standard file system interfaces and semantics that allow you to migrate your enterprise applications to Alibaba Cloud or construct new applications.

- Media and entertainment workflows

You can use Apsara File Storage NAS to share and process large files in media workflows, such as video editing, audio and video production, broadcast processing, and audio design and rendering. Apsara File Storage NAS provides powerful data consistency models, high throughput, and shared access to files. This reduces the time required to complete the preceding workflows and merges multiple on-premises file repositories into a single repository that can be accessed by all users.

- Big data analysis

Apsara File Storage NAS provides high throughput for computing nodes, read and write consistency, and low latency to meet the scale and performance requirements of big data applications. Many analysis workloads use file system interfaces to access data based on file semantics such as file locking, and need to write data to files. In this case, you can use Apsara File Storage NAS that supports file system semantics such as file locking and provides scalable capacity and performance.

Performance

The peak throughput of a file system is linearly proportional to the used capacity of the file system. A file system with larger capacity has higher peak throughput. Apsara File Storage NAS can be concurrently accessed and randomly read or written by thousands of ECS instances by using POSIX.

Specifications	Capacity	Latency	IOPS
Extreme NAS	256 TB	About 100 microseconds	10,000 to 200,000
NAS Performance	1 PB	Milliseconds	Up to 30,000 (4K random read/write)
NAS Capacity	10 PB	About 10 milliseconds	Up to 15,000 (4K random read/write)

Operations

Apsara File Storage NAS operations on data, such as read and write operations, are implemented by using POSIX. You can easily migrate local applications to the cloud without modifying the code.

Apsara File Storage NAS management operations can be used to send GET and POST requests by using HTTP or HTTPS. If you are familiar with network protocols and one or more programming languages, we recommend that you call API operations to manage your Apsara File Storage NAS resources. You can use Apsara File Storage NAS SDKs, Alibaba Cloud Command Line Interface (CLI), or Alibaba Cloud API Explorer to call Apsara File Storage NAS management operations to perform the following operations on file systems, mount targets, permission groups, snapshots, and tags: create, delete, query, and modify. For more information, see [API operations](#). If you prefer graphical web application, you can use the [Apsara File Storage NAS console](#) to perform all operations supported by Apsara File Storage NAS management operations.

Scalability and elasticity

When you use Apsara File Storage NAS, you do not need to perform complex operations that you perform on traditional storage systems, such as planning, purchasing, partitioning, and monitoring. The capacity of an Apsara File Storage NAS file system automatically scales in or out when you delete files from or add files to the file system. This way, Apsara File Storage NAS allocates storage resources based on your requirements without affecting your applications.

Data durability and service availability

Apsara File Storage NAS provides multiple replicas for each piece of data that is stored in a file system. These replicas reside in devices that are isolated across different fault domains for geo-redundancy. Apsara File Storage NAS provides data reliability of 99.999999999% (eleven 9's). This reduces a large number of data security risks.

Security

- Permission group

In Apsara File Storage NAS, a permission group is a whitelist that defines the permission information of a file system, including the authorized IP addresses, read and write permissions, and the permissions of users. You can add rules to a permission group to allow access to a file system from specific IP addresses or CIDR blocks. You can also grant different access permissions to different IP addresses or CIDR blocks.

- RAM

You can use Resource Access Management (RAM) to manage the users of Apsara File Storage NAS and control the access permissions of resources. RAM implements access control based on users. You can create and manage multiple RAM users under an Alibaba Cloud account. You can also grant different permissions to each RAM user. This allows each RAM user to have different access permissions on Alibaba Cloud resources. By using RAM, you do not need to share your AccessKey pairs with other users. You can assign minimal permissions to each RAM user to reduce data security risks for your enterprise. For more information, see [Perform access control based on RAM policies](#).

- ACL

You can use Access Control Lists (ACLs) to manage the access permission of files and directories. ACL implements access control based on resources. Access control and user management are necessary for enterprises that want to share files among different users and groups by using a shared file system. Apsara File Storage NAS provides the ACL feature that allows you to grant users and groups different access permissions on directories and files. For more information, see [NAS NFS ACL](#) and [NAS SMB ACL](#).

- Encryption

Apsara File Storage NAS uses the 256-bit advanced encryption standard (AES-256) to encrypt static data stored in file systems and uses Key Management Service (KMS) to manage encryption keys. Apsara File Storage NAS encrypts data before it writes the data to file systems and automatically decrypts encrypted data before it reads the data and provides the data to applications. Apsara File Storage NAS automatically encrypts and decrypts data. Therefore, you do not need to modify your application code for data encryption and decryption.

Cost model

The capacity of an Apsara File Storage NAS file system automatically scales in or out based on your business requirements. Therefore, you do not need to partition the file system in advance. Accordingly, Apsara File Storage NAS is billed in the pay-as-you-go mode. You are charged only for the storage space that is actually used. If a file is deleted from a file system, the storage space that was used to store the file no longer incurs fees. You can also purchase storage plans in advance to deduct the resource that you use. In most cases, storage plans are more cost-effective.

3.4. Block Storage

Block Storage is a high-performance, low-latency block storage service for Alibaba Cloud Elastic Compute Service (ECS). Block Storage is similar to a physical disk. You can format a Block Storage device and create a file system on it to meet the data storage requirements of your business.

Scenarios

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. Examples:

- Cloud disks are block-level storage devices provided by Alibaba Cloud for ECS instances. Cloud disks use a triplicate distributed mechanism and provide 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 provide low latency, high random IOPS and throughput, and high cost performance.

Performance

The key metrics used to measure Block Storage performance include IOPS, throughput, and latency. Some Block Storage devices also have requirements on capacity. For example, enhanced SSDs of different performance levels (PLs) have different capacity ranges. For more information, see [EBS performance](#).

Data durability and service availability

Three copies of your business data are stored in the Block Storage cluster in the same zone to ensure 99.9999999% (nine 9's) data reliability during read and write operations.

To improve the reliability of Block Service, we recommend that you create snapshots on a regular basis to provide data backup capabilities for cloud disks to make sure that information such as logs and customer transactions are backed up. For more information, see [Snapshot overview](#).

Scalability and elasticity

You can partition or release a cloud disk and adjust the storage capacity in real time based on your business requirements. You can scale out the cloud disk attached to your ECS instance to meet the storage requirements by using the following methods when your business and application data grow:

- Resize the disk. In this case, you must resize the existing partitions of the disk or create new partitions for the disk.
- Create a disk, attach the disk to an ECS instance, and then partition and format the disk.
- Replace the system disk of an instance and specify a greater size for the new system disk.

Security

You can use Resource Access Management (RAM) to authorize other users to access your cloud disks.

We recommend that you encrypt the storage devices that you use if your applications are data-sensitive. Cloud disks and their snapshots are encrypted by using keys based on the standard AES-256 algorithm. Data is automatically encrypted when it is transmitted from ECS instances to cloud disks. Encrypted data is automatically decrypted when it is read. For more information, see [Encryption overview](#).

API operations

Block Storage provides API operations that you can use to send GET and POST requests by using HTTP or HTTPS. If you are familiar with network protocols and one or more programming languages, we recommend that you call API operations to manage your Block Storage resources. You can use ECS SDKs, Alibaba Cloud Command Line Interface (CLI), or Alibaba Cloud API Explorer to call Block Storage API operations to perform the following operations: create, delete, query, attach, detach, and scale out cloud disks on ECS instances and create, delete, and query the cloud disk snapshots stored in OSS.

If you prefer graphical web application, you can use the ECS console to perform all operations supported by Block Storage API operations.

Cost model

Block Storage supports two billing methods: pay-as-you-go and subscription. If you use the pay-as-you-go method, you can activate and release Block Storage resources based on your requirements and do not need to purchase a large number of resources in advance. This way, you can reduce the costs by 30% to 80% compared with traditional hosts. If you use the subscription mode to purchase resources in advance, you can further reduce the costs.

The billing methods of cloud disks depend on how the disks are created.

- Cloud disks created along with an ECS instance use the same billing method as the ECS instance.
- Cloud disks created for a subscription ECS instance are billed by using the subscription method.
- Cloud disks created on the Disks page of the ECS console support only the pay-as-you-go method.
- Cloud disks created from snapshots support only the pay-as-you-go method.

You can change the billing methods of your cloud disks. For more information, see [Change the billing methods of a disk](#). SCUs can be used to offset bills of eligible pay-as-you-go cloud disks in the current

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

3.5. Cloud Paralleled File System

Cloud Paralleled File System (CPFS) is a parallel file system provided by Alibaba Cloud. CPFS stores data across multiple data nodes in a cluster and allows data to be simultaneously accessed by multiple clients. Therefore, CPFS can provide data storage services with high input/output operations per second (IOPS), high throughput, and low latency for large-sized, high-performance computing clusters.

Scenarios

CPFS is optimized for high-performance computing scenarios, and enables users to access data in milliseconds and read and write data with highly aggregated I/O and high IOPS. It can be used in a variety of scenarios such as AI deep training, autonomous driving, genomics computing, EDA simulation, petroleum exploration, meteorological analysis, machine learning, big data analytics, and film rendering.

Performance

CPFS can provide a bandwidth of hundreds of Gbit/s, an IOPS of millions, and sub-millisecond level latency. The bandwidth and IOPS vary with the specification of the file system that you purchase.

Capacity	Bandwidth	Write IOPS of Standard Edition	Read IOPS of Standard Edition	Write IOPS of Advanced Edition	Read IOPS of Advanced Edition
2 TB	1 Gbit/s	20,000	40 k	50 k	100 k
5 TB	1 Gbit/s	20,000	40,000	50,000	100,000
	2.5 Gbit/s	60,000	150,000	100,000	350,000
10 TB	1.5 Gbit/s	40,000	100,000	80,000	200,000
	2.5 Gbit/s	60,000	150,000	100,000	350,000
30 TB	2.5 Gbit/s	60,000	150,000	100,000	350,000
	5 Gbit/s	80,000	200,000	140,000	500,000
50 TB	2.5 Gbit/s	60,000	150,000	100,000	350,000
	7.5 Gbit/s	100,000	300,000	200,000	600,000

You can use the FIO tool to test the throughput and IOPS of a CPFS file system, and configure the number and size of stripes to improve the aggregated bandwidth and IOPS of the file system. For more information, see [Performance](#) and [Performance optimization](#).

Data durability and service availability

CPFS persistently stores data in the Apsara Distributed File System developed by Alibaba Cloud. CPFS supports multiple data copies and provides 99.99999999% (eleven 9's) data reliability.

All nodes in CPFS are designed in the highly available mode. This allows CPFS to detect faults in the cluster within a few seconds and use the cluster scheduler to automatically switch services to other nodes with load balancing at the same time. The failover process is imperceptible to users. Therefore, CPFS provides much higher availability than the traditional two-node mode.

Scalability and elasticity

CPFS supports online scale-out. CPFS stores all data as stripes and supports automatic load balancing after scale-out. This allows you to improve the performance of file systems linearly and use the throughput and storage capabilities of added nodes immediately to meet the growing requirements on capacity and performance.

Security

CPFS allows you to control access to file systems by using an enterprise-created Lightweight Directory Access Protocol (LDAP) service. If LDAP is not integrated, CPFS allows only the root user to access file systems. If other users attempt to access the file systems, the permission denied error is returned. If you use LDAP to access CPFS, you must specify the parameters of the LDAP server and make sure that the LDAP service is available.

Operations

CPFS provides API operations that you can use to send GET requests by using HTTP or HTTPS. If you are familiar with network protocols and one or more programming languages, we recommend that you call API operations provided by CPFS to create, mount, and manage file systems and manage LDAP users.

If you prefer graphical web application, you can use the console to manage CPFS file systems.

Cost model

CPFS are billed based on the storage capacity and bandwidth of file systems. By default, CPFS is billed based on the used resources on an hourly basis (pay-as-you-go). You can also purchase resource plans (subscription) in advance to obtain additional discounts. For more information, see [Billable item and billing methods](#).

3.6. Apsara File Storage for HDFS

Apsara File Storage for HDFS is a file storage service for computing resources such as Alibaba Cloud ECS instances and Container Service. You can manage and access data in Apsara File storage for HDFS in the same way as you do in Hadoop distributed file systems. You can use Apsara File Storage for HDFS without modifying your big data analytics applications. Apsara File Storage for HDFS provides unlimited capacity, performance scaling, unique namespace, multi-party sharing, high reliability, and high availability.

Scenarios

[Apsara file storage for HDFS](#) is suitable for scenarios where high throughput is required, such as big data analytics and machine learning. Apsara File Storage for HDFS supports high-throughput and low-latency access. You do not need to migrate data to local computing resources.

After data is stored in Apsara File Storage for HDFS, ECS instances or other computing resources can directly access the data. You can deploy Hadoop or other machine learning applications on multiple computing resources so that applications can access data in Apsara File Storage for HDFS by using the Hadoop's operation to perform online or offline computing. You can also export the computing results to Apsara File Storage for HDFS to permanently store the results.

Performance

The performance of Apsara File Storage for HDFS is measured by throughput. The practical throughput of an Apsara File Storage for HDFS file system cannot exceed the maximum bandwidth of the ECS instance to which the file system is attached. For example, if the bandwidth of the ECS instance is 1.5 Gbit/s, the throughput of the file system can reach a maximum of 187.5 Mbit/s. The throughput of an Apsara file storage for HDFS file system depends on the capacity of the file system.

Data durability and service availability

Like Apsara File Storage NAS, Apsara File Storage for HDFS provides multiple replicas for each piece of data that is stored in a file system. These replicas reside in devices that are isolated across different fault domains for geo-redundancy. Apsara File Storage for HDFS provides data reliability of 99.999999999% (eleven 9's). This reduces a large number of data security risks.

Scalability and elasticity

Apsara File Storage for HDFS provides your applications with optimal storage performance, including high throughput, high IOPS, and low latency. Additionally, the linear relationship between performance and capacity can meet your requirements on capacity and storage performance when your business grows.

Security

Apsara File Storage for HDFS uses multiple security mechanisms to guarantee the security of data stored in file systems. These security mechanisms include network isolation based on VPCs, user isolation in classic networks, standard permission control for file systems, access control based on security groups, and RAM user authorization.

Operations

Apsara File Storage for HDFS provides SDKs for file systems and management systems. Only SDKs for file systems are provided during the public preview. You can perform management operations in the Apsara File Storage for HDFS console. Apsara File Storage for HDFS SDK for Java implements operations based on Hadoop distributed file systems to provide Hadoop-compatible file systems. The SDK is provided as a JAR file whose name is in the following format: `aliyun-sdk-dfs-x.y.z.jar`. Computing and analysis applications based on Apache Hadoop, such as MapReduce, Hive, Spark, and Flink, can use the SDK to use Apsara File Storage for HDFS as the default file system without modifying and compiling code for better features and performance than the original HDFS.

If you prefer graphical web applications, you can use the console to manage Apsara File Storage for HDFS file systems.

Cost model

Apsara File Storage for HDFS is billed based on the capacity and preset throughput of the file system. By default, Apsara File Storage for HDFS is billed based on the used resources on an hourly basis (pay-as-you-go). You can also purchase resource plans (subscription) in advance to obtain additional discounts. For more information, see [Pricing](#).

3.7. Tablestore

Tablestore is a data storage service developed by Alibaba Cloud to store a large amount of structured data. You can use Tablestore to efficiently query and analyze data. Tablestore provides the Wide Column model, Timeline model, and Timestream model that are compatible with HBase. It can store petabytes of data and provide tens of millions of transactions per second (TPS) and milliseconds of latency.

Scenarios

Tablestore can store petabytes of data in a single table. In addition, it supports tens of millions of query per second (QPS) and a variety of query methods, including global secondary index, full-text search, inverted index, and spatio-temporal index. Therefore, Tablestore is widely used in scenarios of structured data, such as social networks, Internet of Things (IoT), artificial intelligence (AI), metadata, and big data.

- Metadata

When you store a large amount of data such as documents and media files, it is important to store and analyze the metadata of the stored data. In scenarios where e-commerce orders, transaction histories of bank accounts, and phone bills from service providers are stored, you also need to store and analyze a large amount of metadata. Tablestore can help you efficiently manage the metadata of your data.

- Message data

Tablestore provides the Timeline model to store message data. This model constructs lightweight message queues that support a large number of topics to store a large amount of information in social network applications, such as instant messaging (IM) chats and Feed streams information such as comments, posts, and likes. The Timeline model is already used in a variety of IM systems, such as DingTalk, to support the synchronization of a large number of chat messages.

- Trajectory tracing

Tablestore provides the Timestream model to help you manage and analyze trajectory data in different scenarios, such as running, riding, walking, and food delivery.

- Scientific big data

Gridded data is a type of scientific big data used in geoscience fields such as meteorology, oceanography, geology, or geomorphology. Gridded data is widely used and the data volumes are growing. Scientists need to quickly browse data and query data online by using various methods, which requires a low latency. Tablestore can meet the high requirements on storage capacity and query performance in scientific big data scenarios.

- Internet big data

Product designers of e-commerce and information platforms on the Internet must collect and analyze the data of different platforms to determine the further development of their products. The public relations and marketing departments of enterprises also need to handle issues in a timely manner based on public opinions. Tablestore can help you store and analyze tens of billions of public opinions.

- IoT

Tablestore can be used to store time series data from IoT devices and monitoring systems. It provides API operations to directly read SQL data and incremental data streams, which allow you to implement offline data analysis and real-time stream computing.

Performance

Tablestore can store tens of petabytes of data and trillions of records in a single table and provide tens of millions of transactions per second (TPS) and milliseconds of latency. It supports automatic load balancing and hotspot migration without manual operations and maintenance (O&M). In addition, Tablestore can provide high throughput for write operations and stable read and write performance that can be predicted. For more information, see the [Tablestore performance white paper](#).

Data durability and service availability

Tablestore creates multiple backups of data and stores them in different servers across racks. When a backup fails, Tablestore immediately uses another backup to restore the data. This mechanism ensures data durability of 99.99% and service availability of 99.99999999% (eleven 9's).

Scalability and elasticity

Tablestore uses shards and load balancing to implement seamless scalability. Tablestore adjusts the size of partitions to store more data in a table when the data stored in the table increases. Tablestore can store a minimum of 10 PB of data. A single table in Tablestore can store a minimum of 1 PB of data or one trillion records.

Security

Tablestore performs authentication and authorization based on tables and operations. It also supports STS temporary authorization, custom authentication, and RAM users to isolate resources by user. For more information, see [RAM and STS](#). Tablestore supports access from the Internet, ECS instances, and VPCs and provides network access control.

Operations

Tablestore provides standard RESTful API operations. You can use tools or Tablestore SDKs that encapsulate the API operations to develop applications. Tablestore provides SDKs for [various programming languages](#), including Java, Python, PHP, and Go. [TableStore CLI](#) is a command-line interface (CLI) tool that provides simple and clear commands to perform a variety of operations in Tablestore, including table operations, single-row operations, simple stress testing operations, and data backup operations. This tool supports Windows, Linux, and Mac operating systems.

The Tablestore console allows you to create instances, tables, and search indexes, perform basic read and write operations on data, and monitor the access data, such as QPS, latency, and number of requests, of instances and tables.

Cost model

You can use the pay-as-you-go billing method to pay only for the Tablestore resources that you use. This way, you can handle traffic fluctuations and high concurrency requests that require low latency at low costs. You can also use the subscription billing method to purchase resource plans in advance to deduct fees incurred by resource usage. Tablestore is a pay-as-you-go service that is billed based on the following items: storage usage, read throughput, write throughput, and Internet outbound traffic. If you use the search index or global secondary index feature, additional fees are incurred. For more information, see [Billing overview](#).

3.8. Cloud Storage Gateway

Cloud Storage Gateway (CSG) is a gateway service that can be deployed in on-premises data centers and Alibaba Cloud. CSG uses Alibaba Cloud Object Storage Service (OSS) buckets as backend storage devices and provides on-premises and in-cloud applications with standard file services by using the Network File System (NFS) and Server Message Block (SMB) protocols and block storage services by using the Internet Small Computer Systems Interface (iSCSI) protocol.

CSG supports two types of gateways:

- File gateways

File gateways map the objects and folders in OSS buckets to the files and directories in Apsara File Storage NAS file systems. This way, you can read and write all objects in a specified OSS bucket by using standard NFS and SMB protocols. File gateways use local storage as the cache for hot data to provide high-performance data access to the large storage space provided by OSS buckets. File gateways are compatible with Portable Operating System Interface (POSIX) and third-party backup software. If you want to back up small files or share small files for reading and writing, we recommend that you use standard or basic types of file gateways. If you have high performance requirements or want to use multiple clients to simultaneously access shared data, we recommend that you use enhanced and advanced types of file gateways.

- Block gateways

Block gateways create storage volumes in OSS and allows you access OSS by using the Internet Small Computer Systems Interface (iSCSI) protocol. Local applications can access these volumes as iSCSI targets. Block gateways run in two modes: write-through and cache. In the write-through mode, data stored in volumes is sliced and synchronized to the cloud. This mode is applicable to scenarios where you use high-speed links such as leased lines. In the cache mode, local cache disks are created to accelerate read and write operations and transfer the cached data to the cloud asynchronously. This mode is applicable to scenarios where you want to access locally cached data quickly and transfer data to OSS in a normal speed.

Scenarios

CSG is applicable to the following scenarios:

- Storage scale-out and data migration

CSG is integrated with intelligent caching algorithms, which automatically identify hot and cold data. This allows CSG to store cold data in the cloud and retain hot data in the local cache to accelerate access to hot data. When you connect your on-premises data center to cloud storage services, your users can access the requested data without noticing the connection. This enables you to expand your storage from your own data center to cloud storage. CSG also stores a full copy of data in the cloud to guarantee data integrity.

- Cloud disaster recovery

With the development of cloud computing, an increasing number of users run their workloads in the cloud. Therefore, the reliability and continuity of workloads running in the cloud become critical issues. CSG supports virtualization technologies. You can easily connect third-party services to Alibaba Cloud services to implement disaster recovery.

- Data sharing and distribution in multiple regions

You can deploy CSG instances in multiple regions and associate the CSG instances with the same OSS bucket to quickly share and distribute files across multiple regions. This feature is applicable to branch offices where data needs to be synchronized and shared.

- Compatibility with legacy applications

Some users may run both legacy and modern applications in the cloud. In this case, the legacy applications migrated from an on-premises data center use standard storage protocols, such as NFS, SMB, and iSCSI. Modern applications are typically developed based on new technologies and support object access protocols. Data communication among applications that use different protocols requires complicated processes. CSG can establish communications among applications that use different protocols and enable data exchange between legacy and modern applications.

- An alternative to ossfs and ossftp

ossfs and ossftp are open-source tools developed based on file protocols. You can use ossfs and ossftp to upload files to OSS. However, ossfs and ossftp are not supported in the production environment due to their low compatibility with POSIX. If you use ossfs and ossftp to mount file systems to a client, additional configurations and caches are required. In scenarios where you need to use ossfs and ossftp to mount file systems to multiple clients, the configuration process may take a long time. CSG can be used as an alternative to ossfs and ossftp. To accelerate access to data stored on OSS, you need only to create a CSG instance and mount an NFS share to your local client in Linux or map an SMB share to a network drive in Windows. You can then manage data in the OSS bucket in the same way as you manage it in the local file system.

Performance

CSG is deployed between applications and OSS. Therefore, the performance of CSG is determined by multiple factors, including the speed and configurations of local disks, the bandwidth between the iSCSI initiator and the gateway, the storage capacity allocated to the gateway virtual machine, and the bandwidth between the gateway virtual machine and OSS. To allow local applications to access data with low latency, you must configure sufficient local cache to store the data that is accessed recently. The cache capacity of a share can be calculated in the following formula: Recommended local cache capacity = [Application bandwidth (Mbit/s) - Backend bandwidth of the gateway (Mbit/s)] × Write duration (seconds) × 1.2. To obtain better performance when you access data from local clients, you can estimate the total amount of hot data. Compare the total amount of hot data with the recommended local cache capacity, and select the larger value as the capacity of the local cache disk. The synchronization bandwidth of a gateway is determined by the bandwidth of OSS. OSS provides up to 10 Gbit/s of bandwidth for each user. The bandwidth slightly varies among clusters in different regions. For more information, contact the technical support of OSS in each region.

CSG supports the transfer acceleration feature that increases the data transfer rate across regions by using the public bandwidth of gateways. You can enable transfer acceleration when or after you create a share.

Data durability and service availability

In the cache mode, data is written to a disk by using synchronous I/O to prevent data loss due to power failures. CSG ensures the durability and reliability of data in the local cache disks by using Alibaba Cloud disks. The reliability of on-premises gateways depend on the reliability of the backend storage in your virtual environment. We recommend that you use Redundant Array of Independent Disks (RAID) storage or a distributed storage system as local cache disks. CSG refreshes and uploads data in the local cache disks to OSS buckets. Alibaba Cloud OSS guarantees 99.999999999% (twelve 9's) designed data reliability. This ensures data security and reliability when you transfer data from CSG to Alibaba Cloud.

Scalability and elasticity

CSG uses OSS buckets as backend storage devices. Therefore, it inherits the scalability and elasticity of OSS. You may encounter problems if you store a large number of files in the same directory in a common file system. In contrast, the total storage capacity of OSS and the capacity of a single bucket are not limited. You can store an unlimited number of objects in a bucket. OSS stores the copies of your data in different servers in the same region. All copies have the same performance provided by the high-performance infrastructure of Alibaba Cloud.

Security

CSG supports user identity management and resource access control based on Resource Access Management (RAM). RAM implements access control based on users. RAM allows you to create and manage multiple RAM users under an Alibaba Cloud account and grant diverse permissions to each RAM user. This way, you can authorize different RAM users to access different Alibaba Cloud resources. You can use RAM to avoid sharing your AccessKey pairs with other users. You can assign minimal permissions to each RAM user to reduce data security risks for your enterprise. For more information, see [Use RAM to implement account-based access control](#).

CSG provides data encryption feature to encrypt the data transferred from gateways to OSS buckets. You can also enable the server-side encryption feature provided by OSS and configure a CMK. This way, files uploaded from a share to OSS are automatically encrypted on the OSS server by using the specified CMK.

Operations

You can perform the following operations on the CSG console: deploy a gateway virtual machine to an on-premises data center or an ECS instance, configure a file gateway or block gateway, and create cache disks and shares. You can also manage CSG by using API operations. CSG provides API operations that you can use to send GET requests by using HTTP or HTTPS.

Cost model

CSG is billed based on the used resources. Gateways deployed on the cloud and on-premises data centers are billed based on different items. Gateways deployed on the cloud are billed based on the gateway specifications, gateway cache types, and the public network bandwidth. Gateways deployed on the cloud supports the pay-as-you-go and subscription billing methods. Gateways deployed on on-premises data centers run with your own virtual machines and cache resources. Alibaba Cloud only charges for the gateway software license. For more information, see [Billable items and billing methods](#).

4. Optimization of Alibaba Cloud storage services

4.1. 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.

4.2. 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.

4.3. 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.

Standard	Standard storage provides high-durability, high-availability, and high-performance object storage services that can handle frequent data access. Standard storage is ideal for storing images for social networking and sharing, storing data for audio and video applications, large websites, and big data analytics.
Infrequent Access (IA)	IA storage is suitable for data that is infrequently accessed, such as only once or twice a month. IA storage offers a storage unit price lower than that of Standard storage and is suitable for long-term data backup of mobile apps, smart devices, and enterprises. It also supports real-time data access.
Archive	Archive storage is suitable for long-term (at least six months) storage of data that is infrequently accessed. OSS takes a minute to restore an Archive object before you can read the object. Archive storage is suitable for data that you want to store for a long period of time such as archival data, medical images, scientific materials, and video footage.
Cold Archive	Cold Archive storage is suitable for extremely cold data that you want to store for an extremely long period of time. Examples: data that must be retained for an extended period of time due to compliance requirements, raw data that is accumulated over an extended period of time in the big data and AI fields, media resources that are retained in the film and television industries, and archived videos from the online education industry.

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.

Extreme NAS	Extreme NAS is a high-performance file sharing solution that is built on top of the latest generation of network architecture and all-flash storage. The maximum capacity is 256 TB. The bandwidth ranges from 150 Mbit/s to 1,200 Mbit/s. Extreme NAS can provide a constant latency of about 100 microseconds. Extreme NAS is applicable to latency-sensitive business in which a large amount of small files are handled.
NAS Performance	NAS Performance uses solid-state drives (SSDs) as the storage devices, and provides high throughput, high input/output operations per second (IOPS), and low latency for workloads. NAS Performance is a file sharing solution that is applicable to scenarios where high throughput, high concurrency, business scalability, and low read latency are required. You can use NAS Performance if you need to perform frequent read/write operations and have high requirements for response latency.

NAS Capacity	NAS Capacity uses SATA hard disk drives (SATA HDDs) as storage devices and provides high-performance storage space at low costs. NAS Capacity is a file sharing solution that is applicable to cost-sensitive scenarios where high throughput, high concurrency, and business scalability are required. NAS Capacity is more cost-efficient if you do not need to perform frequent read/write operations and do not have high requirements on response latency.
Infrequent Access	You can configure lifecycle rules to transfer data that is infrequently accessed and needs to be stored for long periods from NAS Performance or NAS Capacity to Infrequent Access to reduce costs. For more information, see Implementation of lifecycle management .

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.

4.4. 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.

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.

4.5. 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.

4.6. 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.