

Cloud Object Storage

FAQs

Product Documentation



Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by the Tencent corporate group, including its parent, subsidiaries and affiliated companies, as the case may be. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

FAQs

General

Billing

Bucket Configuration

 Lifecycle

 Batch Operations

 Origin-Pull

 Bucket Tags

 Hotlink Protection

 Cross-Origin Resource Sharing

 Inventory

 Deleting Buckets

Domain Names and CDN

 Static Website

 Private Network Domain Name

 Global Acceleration Endpoint

 Custom Origin Domain Name

 CDN Acceleration Domain Names

 Other Domain Name Issues

Object Operations

 Upload and Download

 Data Management

 Storage Class

Logging and Monitoring

Permission Management

Data Processing

 Function Service

Data Security

SDKs

Tools

 COSBrowser

 COS Migration

 COSCMD

 FTP Server

 Hadoop

APIs

FAQs

General

Last updated : 2025-03-13 12:40:10

What is COS?

Tencent Cloud Cloud Object Storage (COS) is a cloud-based non-hierarchical distributed storage service that provides cost-effective, fast, and reliable data storage solutions. COS stores your data across multiple AZs, incorporating redundant storage to ensure data reliability, and allows multiple clients or application threads to read or write data at the same time.

You can use web APIs to store and retrieve data through CVM instances or over the internet. You can also use the URL of a specified domain name to store and retrieve individual data object in COS through HTTP or HTTPS protocol.

For more information about COS, please see [COS Documentation](#).

What is the difference between Cloud Object Storage (COS) and Cloud File Storage (CFS)?

[COS](#) has no limit on directory hierarchy or data format and can store any amounts of data. There is no upper limit on the storage capacity of buckets where no partitioning is required. It supports HA deployment to ensure the eventual consistency of data but not features such as file locking. Its APIs are provided for data access using the HTTP or HTTPS protocol, and its SDKs and tools can be integrated into your businesses. Objects uploaded to COS can be accessed or downloaded directly through URL.

[Cloud File Storage \(CFS\)](#) uses common network file transfer protocols, can create file systems and implement large-scale expansion, but needs to be mounted onto CVM. It can store data for a wide range of applications such as websites, online distribution, and archiving. Featuring high computing throughput and extremely high availability and persistence, it is also suitable for scenarios demanding high concurrence or shared storage.

What is the difference between COS and CBS?

[COS](#) has no limits on file systems, directory structure, number of files, and storage capacity. It needs to be managed and accessed via web APIs. It offers various SDKs and tools for integration, which can also be used separately without CVM. COS supports access to massive amounts of data but is not suitable for scenarios involving millisecond-level response or random I/O.

[Cloud Block Storage \(CBS\)](#) needs to be used together with CVM and can only be mounted and used after the file system is partitioned or formatted. It comes in different types with various performance metrics such as IOPS and throughput for different scenarios.

Why does the access link of a public-read file expire?

If you use an access link with a temporary signature, the link expires when the temporary signature expires, regardless of whether the file is public-read.

If you want your public-read file to be always accessible, we recommend you use an unsigned access link, which is the object URL on the object details page of the [COS console](#).

What is a "folder" or "directory" in COS?

The concepts of folder and directory do not apply to COS. However, taking into account the usage habits of different users, COS displays “folders” in the console and COSBrowser just like in the directory structure of traditional file management. For more information, see [Folder and directory](#).

Can COS files be recovered after being deleted?

The data redundancy storage mechanism of COS is designed for scenarios where it is necessary to recover data in case of hardware failure. When versioning is not enabled, if you manually delete your data from COS or configure automated deletion, Tencent Cloud will delete the data as requested after which the data is irrecoverable.

You can proactively delete files in the following ways:

Deleting files using the COS console by deleting a single file, deleting files in batches, clearing incomplete multipart uploads, or emptying buckets.

Deleting files using COS tools such as COSCMD or COSBrowser.

Deleting files using the COS APIs or SDKs.

Configuring the system to delete files regularly through the COS lifecycle management feature.

Syncing the CRUD operations between buckets in different regions using the COS cross-region replication sync feature, so that existing files with the same name will be overwritten or deleted.

How can I avoid accidental deletion?

The best way to avoid accidental deletion is to back up the files in your bucket on a regular basis. You can protect your data in the following ways:

Download the objects in COS to your local file system or third-party servers using the [COSCMD tool](#).

Perform intra-region or cross-region bucket data backup using the [COS Migration Tool](#) or the cross-region replication feature.

Regularly backing up your data to other COS buckets using COS APIs or SDKs.

Saving your past versions of data using versioning.

Use COS permission management. For more information, see [Cloud Access Management Practices](#):

Separating read permission and write permission. For businesses where it is only necessary to read data, use a sub-account with read permission or a temporary key to access the data.

Separating permissions for different buckets. For different businesses, only authorize permissions for buckets, directories, and operations within the scope of that particular business.

Not using a root account to access COS.

Accessing COS using a temporary key.

Properly managing your data access credentials, such as your Tencent Cloud account password, CAM sub-account access credentials, and TencentCloud API key.

Does COS support statistics collection?

COS is capable of monitoring stored data and displaying the details and trends of various metrics in the monitoring window. To view general data trends, go to the **Overview** page in the [COS console](#), and you can view data such as storage size, request number, and traffic for each storage class.

To view statistics of a single bucket, see [Querying Monitoring Reports](#).

In addition to the COS Console, you can also view the monitoring information of different buckets on the [Cloud Monitoring](#) page where you can also configure different alarm policies to fit your business needs.

Does COS support image processing, image compression, thumbnail generation, or video transcoding?

CI is integrated in the COS console to implement data processing features such as image processing, image compression, thumbnail generation, and video transcoding. For more information, see [Data Processing](#).

What formats of audio/video files can COS process?

COS is a distributed storage service for unstructured data and cannot support image compression or audio/video file processing on its own. For more information on rich media file processing (MP4, AVI, TS, HLS, MP3, AAC, etc.), see [Cloud Infinite](#).

Does COS support the auto decompression of uploaded files?

COS is a distributed storage service for unstructured data and does not support file decompression; however, you can use the SCF service to decompress files. For more information, see [Setting File Decompression](#).

What are the specifications and limits of COS?

For more information, see [Specifications and Limits](#).

What is a bucket?

A bucket is a carrier of objects, which can be considered as a "container" for storing objects. You can manage buckets and configure attributes of buckets in various methods such as the Tencent Cloud console, APIs, and SDKs. For example, you can set a bucket for hosting a static website or set access permission on a bucket. For more information on buckets, see [Bucket Overview](#).

What is the length limit on a bucket name?

The bucket length limit has been changed since the COS console update in September 2021. According to the new policy, the length of a bucket name is affected by the number of characters in the **region abbreviation** and **APPID**,

as the combined full domain can contain 60 characters at most. Note that previous bucket names will not be affected. If you need longer names to meet special requirements, [contact us](#).

How do I monitor error code information?

You can use [Cloud Monitor](#) to get different types of HTTP error code messages. For more information, see [Monitoring and Alarm](#). For information about how to work with Cloud Monitor and obtain relevant data, see Cloud Monitor's [Console Guide](#) or [API Documentation](#).

How do I calculate the availability of COS?

Please reference the following example for information on how to calculate COS availability:

Tom uses Tencent Cloud COS to run his e-commerce business. Assume that his business consumed a total of 100 USD in the service period from Nov. 1 to Nov. 30, 2018, during which two unavailability events occurred, as shown below:

| Unavailability Event No. | Duration | 5-Minute Record of Unavailability Event | HTTP Return Code | Number of Failing Requests | Number of Valid Requests |
|--------------------------|----------|---|------------------|----------------------------|--------------------------|
| 1 | 15 min | November 15, 2018, 10:00 - 10:05 | 503 | 100 | 100 |
| | | November 15, 2018, 10:05 - 10:10 | 503 | 99 | 100 |
| | | November 15, 2018, 10:10 - 10:15 | 503 | 98 | 100 |
| 2 | 15 min | November 20, 2018, 16:00 - 16:05 | 500 | 150 | 150 |
| | | November 20, 2018, 16:05 - 16:10 | 500 | 148 | 150 |
| | | November 20, 2018, 16:10 - 16:15 | 500 | 140 | 150 |

In all other periods, Tom's requests **were successful and a 200 status code was returned**.

In this case, the overall availability for the service period is as follows:

(1) Calculate the per-5-minute error rate for the current month

According to the case details: when Tom's business is normal, the per-5-minute error rate is 0%.

Unavailability event 1: This event occurred on November 15, 2018 and lasted from 10:00 - 10:15. The per-5-minute error rate was:

$100 / 100 * 100\% = 100\%$ from 10:00 - 10:05

$99 / 100 * 100\% = 99\%$ from 10:05 - 10:10

$98 / 100 * 100\% = 98\%$ from 10:10 - 10:15

Unavailability event 2: This event occurred on November 20, 2018 and lasted from 16:00 - 16:15. The 5-minute error rate was:

$150 / 150 * 100\% = 100\%$ from 16:00 - 16:05

$148 / 150 * 100\% = 98.67\%$ from 16:05 - 16:10

$140 / 150 * 100\% = 93.33\%$ from 16:10 - 16:15

(2) Calculate the service availability for the current month

In this case:

Total duration of the service period: 30 days * 24 hours/day * 60 minutes/hour=43,200 minutes.

Total number of 5-minute periods: 43,200 minutes / 5 minutes = 8,640.

Total number of unavailable 5-minute periods: (15 + 15) minutes / 5 minutes = 6.

Sum of the per-5-minute error rates: (100% + 99% + 98% + 100% + 98.67% + 93.33%) + (8640 - 6) * 0% = 589%

The service availability for this month: $(1 - 589\% / 8640) * 100\% = 99.93\%$

(3) Calculate relevant compensation

In this example, the service availability is 99.93%, which is lower than the standard 99.95% but higher than 99.9%.

According to compensation standards, Tom is eligible for compensation equivalent to 20% of the total monthly service fees, i.e., 20 USD.

Tom only needs to submit a ticket to apply for compensation within sixty (60) calendar days after the end of the service period, i.e., prior to January 29, 2019, and Tencent Cloud will compensate Tom for his losses by issuing a voucher.

How do I deactivate the COS service and stop being charged?

You can deactivate COS or stop its billing as follows:

1. If you decide to stop using COS, you can avoid any further billing by permanently deleting all of your COS data (including incomplete multipart uploads and object versions) as instructed in [Payment Overdue](#). There is no need to remove your account, and if you use other Tencent Cloud products, avoid doing so as it will affect your other services.
2. If you don't use COS for more than one month, you can set lifecycle rules to transition data in STANDARD storage class in the bucket to a colder class such as STANDARD_IA, ARCHIVE, or DEEP ARCHIVE to reduce storage fees. For more information, see [Setting Lifecycle](#). The transition will generate read requests in the original storage class and write requests in the target storage class, so transition by lifecycle will incur read/write [request fees](#).

Note

Data, once deleted from the bucket, cannot be recovered, so make backups accordingly.

If versioning is enabled for the bucket, disable it before deleting data.

Check your billing cycle to avoid overdue payments. If all your billable items are settled daily, the bill for the day of data deletion will be generated on the next day. After the data is completely cleared, the system will stop billing. For more information, see [Billing Cycle](#).

If your account has overdue payment (i.e., your account balance is below 0), COS services will be suspended after 24 hours, regardless of whether your resource pack is within the validity period.

If your account has overdue payment and COS services are suspended, the free tier for which your account is eligible won't be available.

If data in your bucket is blocked for the second time due to non-compliance, it cannot be deleted. [Contact us](#) if you have any questions.

Billing

Last updated : 2025-04-10 19:52:06

Billing Mode

How is COS billed?

COS is pay-as-you-go. For more information, see [Billing Overview](#).

What are the billable items of COS?

COS billable items include [storage usage](#), [requests](#), [data retrievals](#), [traffic](#), and [management features](#). For more information, see [Pricing | Cloud Object Storage](#).

Can COS be billed by bandwidth?

No. COS can be billed only on a pay-as-you-go (postpaid) basis.

How are COS request fees calculated?

Request fees are calculated based on the number of requests sent to COS, including **user requests** and **backend requests** generated after you configure a feature. For more information, see [Request Fees](#).

What changes have been made to COS pricing?

On September 30, 2021, the published prices of COS were reduced as follows:

1. Unit price of object tagging

Product pricing: The prices for public cloud regions in and outside the Chinese mainland were reduced to 0.00025817 USD/10,000 tags/day and 0.0003098 USD/10,000 tags/day, respectively.

Billing cycle: Fees incurred between 00:00 and 23:59:59 on a day are settled the next day.

Bill description: These prices have taken effect for bills generated starting from October 1, 2021 (i.e., fees incurred in September until September 30, 2021).

2. Unit price of DEEP ARCHIVE read/write requests

Product pricing: The prices for all public cloud regions were reduced to 0.07 USD/10,000 requests.

Billing cycle: Fees incurred in a month are billed on the first day of the next month.

Bill description: These prices have taken effect for bills generated starting from October 1, 2021 (i.e., fees incurred in September).

Which regions will benefit from the reduced unit pricing of STANDARD IA storage usage?

The reduction in the unit price of STANDARD_IA storage usage from 0.018 USD/GB/month to 0.015 USD/GB/month will apply to certain regions, including Seoul, Bangkok, Silicon Valley, Jakarta, and São Paulo. The reduction has

taken effect on July 1, 2022.

Price Related Content

Which COS Billable Items Have Adjusted Prices?

Details of the price adjustment records for COS billable items are as follows:

| Adjustment Time | Adjustment Type | Adjusted Billing Item | Pricing Details |
|------------------|-------------------------------|-------------------------------|--|
| January 15, 2024 | Pay-as-You-Go Price Reduction | STANDARD storage capacity | <p>Billable item: STANDARD storage capacity in COS</p> <p>Price: The published price for the public cloud regions within the Chinese Mainland has been reduced to 0.016 USD/GB/month. For Hong Kong (China) and overseas regions, the published price has been reduced to 0.017 USD/GB/month.</p> <p>Billing cycle: Daily. Each day, the cost from 00:00:00 to 23:59:59 (UTC+8) of the previous day is billed.</p> <p>Note:</p> <p>(1) The price will take effect on the bill dated January 15, 2024, reflecting resource usage during the period from 00:00:00 to 23:59:59 (UTC+8) on January 14, 2024.</p> <p>(2) If the existing price for a region is lower than the target price after reduction, the price for that region will remain unchanged.</p> |
| January 15, 2024 | Pay-as-You-Go Price Reduction | MAZ_STANDARD storage capacity | <p>Billable item: MAZ_STANDARD storage capacity in COS</p> <p>Price: The published price for the public cloud regions within the Mainland China has been reduced to 0.0195 USD/GB/month. For Hong Kong (China) and overseas regions, the published price has been reduced to 0.002 USD/GB/month.</p> <p>Billing cycle: Daily. Each day, the cost from 00:00:00 to 23:59:59 (UTC+8) of the previous day is billed.</p> |

| | | | |
|------------------|-------------------------------|----------------------------------|--|
| | | | <p>Note:</p> <p>(1) The price will take effect on the bill dated January 15, 2024, reflecting resource usage during the period from 00:00:00 to 23:59:59 (UTC+8) on January 14, 2024.</p> <p>(2) If the existing price for a region is lower than the target price after reduction, the price for that region will remain unchanged.</p> |
| January 15, 2024 | Pay-as-You-Go Price Reduction | STANDARD_IA storage capacity | <p>Billable item: STANDARD_IA storage capacity in COS</p> <p>Price: The published price for the public cloud regions within the Mainland China has been reduced to 0.01 USD/GB/month. For Hong Kong (China) and overseas regions, the published price has been reduced to 0.011 USD/GB/month.</p> <p>Billing cycle: Daily. Each day, the cost from 00:00:00 to 23:59:59 (UTC+8) of the previous day is billed.</p> <p>Note:</p> <p>(1) The price will take effect on the bill dated January 15, 2024, reflecting resource usage during the period from 00:00:00 to 23:59:59 (UTC+8) on January 14, 2024.</p> <p>(2) If the existing price for a region is lower than the target price after reduction, the price for that region will remain unchanged.</p> |
| January 15, 2024 | Pay-as-You-Go Price Reduction | MAZ_STANDARD_IA storage capacity | <p>Billable item: MAZ_STANDARD_IA storage capacity in COS</p> <p>Price: The published price for the public cloud regions within the Mainland China has been reduced to 0.0132 USD/GB/month. For Hong Kong (China) and overseas regions, the published price has been reduced to 0.0138 USD/GB/month.</p> <p>Billing cycle: Daily. Each day, the cost from 00:00:00 to 23:59:59 (UTC+8) of the previous day is billed.</p> <p>Note:</p> <p>(1) The price will take effect on the bill dated January 15, 2024, reflecting resource usage during the period from 00:00:00 to 23:59:59 (UTC+8) on January 14, 2024.</p> <p>(2) If the existing price for a region is lower than</p> |

| | | | |
|-------------------|-------------------------------|---|--|
| | | | the target price after reduction, the price for that region will remain unchanged. |
| January 15, 2024 | Pay-as-You-Go Price Reduction | ARCHIVE storage capacity | <p>Billable item: ARCHIVE storage capacity in COS</p> <p>Price: The published price for the public cloud regions within the Mainland China has been reduced to 0.004 USD/GB/month. For Hong Kong (China) and overseas regions, the published price has been reduced to 0.0045 USD/GB/month.</p> <p>Billing cycle: Daily. Each day, the cost from 00:00:00 to 23:59:59 (UTC+8) of the previous day is billed.</p> <p>Note:</p> <p>(1) The price will take effect on the bill dated January 15, 2024, reflecting resource usage during the period from 00:00:00 to 23:59:59 (UTC+8) on January 14, 2024.</p> <p>(2) If the existing price for a region is lower than the target price after reduction, the price for that region will remain unchanged.</p> |
| January 15, 2024 | Pay-as-You-Go Price Reduction | INTELLIGENT_TIERING storage capacity | <p>Billable item: INTELLIGENT_TIERING storage capacity in COS</p> <p>Price: The published price for the public cloud in Hong Kong (China) and overseas regions is the same as that of the converted storage.</p> <p>Billing cycle: Daily. Each day, the cost from 00:00:00 to 23:59:59 (UTC+8) of the previous day is billed.</p> <p>Note:</p> <p>(1) The price will take effect on the bill dated January 15, 2024, reflecting resource usage during the period from 00:00:00 to 23:59:59 (UTC+8) on January 14, 2024.</p> <p>(2) If the existing price for a region is lower than the target price after reduction, the price for that region will remain unchanged.</p> |
| November 15, 2023 | Pay-as-You-Go Price Reduction | Management Features - Select Feature Fees | <p>Billable items: COS STANDARD Select feature fee, STANDARD_IA Select feature fee, MAZ_STANDARD Select feature fee, and MAZ_STANDARD_IA Select feature fee</p> |

| | | | |
|--|--|--|---|
| | | | <p>Price: The published price for the public cloud in the Frankfurt region has been reduced from 0.0025 USD/GB to 0.00225 USD/GB.</p> <p>Billing cycle: Daily. Each day, the cost from 00:00:00 to 23:59:59 (UTC+8) of the previous day is billed.</p> <p>Note: The price will take effect on the bill dated November 15, 2023, reflecting resource usage during the period from 00:00:00 to 23:59:59 (UTC+8) on November 14, 2023.</p> |
|--|--|--|---|

How Do I View the Latest Prices for COS?

For the latest prices of COS, you can refer to [Product Pricing](#), [Price Calculator](#), as well as [Traffic Fees](#), [Storage Usage Fees](#), [Request Fees](#), [Data Retrieval Fees](#), and [Management Fees](#) on the official website documentation.

How Does the Adjusted Published Prices for COS Pay-As-You-Go (post-payment) Become Effective?

Whether the adjusted published prices for the COS pay-as-you-go billable items take effect depend on the discount settings of your account.

Activation Logic

Discount price: After the prices are adjusted, the customer automatically uses the latest prices.

Contract price:

After adjustment - Customers with contract price configured: The original contract price prevails, and the latest price is not applicable.

After adjustment - Customers without contract price configured: The latest published price prevails, and the latest price is automatically used.

Effective Scenario

Depending on the discount configuration of the adjusted billable items, whether the price after adjustment can be used is determined. Please refer to the following scenarios for details:

| Billable Item Discount Configuration | Description |
|---|--|
| Adjusted Billable Item: No discount, no contract price | As the adjustment to the published price becomes effective, the published price in the billable item automatically refreshes . No action is required. |
| Adjusted Billable Item: No discount, contract price available | As the adjustment to the published price becomes effective, the contract price remains unchanged and is not affected . There is no change in the discounted price of the billable item on the bill. |
| | |

| | |
|---|--|
| Adjusted Billable Item: Discount available, no contract price | As the adjustment to the published price becomes effective, the published price and discounted unit price in the billable item automatically refresh . No action is required. |
| Adjusted Billable Item: Discount available, contract price available | As the adjustment to the published price becomes effective, the contract price remains unchanged and is not affected . There is no change in the discounted price of the billable item on the bill. |

How Does Reduction in the Published Prices for Pay-As-You-Go COS Resource Packs Become Effective?

After the published price of COS resource pack is adjusted, the new price becomes effective immediately.

Whether your resource pack is billed according to the latest price depends on whether the order is placed after the adjustment (new purchases, renewals, or upgrades). Details are as follows:

| Scenario | Description |
|--|--|
| No resource packs under the account | When a new resource pack is purchased, you can use the latest price after the adjustment. |
| Resource packs available under the account | Before the price adjustment, new purchases, upgrades, and renewals of resource packs still adhere to the original price. After the price adjustment, new purchases, upgrades, and renewals of resource packs will use the latest price. |

Free Tier

Does COS offer a free tier?

COS offers a limited free tier to all new users (both individual and enterprise users) to deduct the fees incurred by data stored in the STANDARD storage class. For more information, see [Free Tier](#).

Does COS offer a free tier outside the Chinese mainland?

Yes. The free tier is applicable to **public cloud regions** (including those outside the Chinese mainland). For more information, see [Regions and Access Endpoints](#).

The free tier offered by COS can be used to deduct only **STANDARD storage usage** but not **other** billable items, such as STANDARD_IA storage usage, ARCHIVE storage usage, requests, and traffic. For more information, see [Billable Items](#).

Why is my account overdue or charged even if I am on the free tier?

The following are some of the possible reasons why your account has overdue payments if your free tier hasn't expired:

Multiple billable items are used, but the free tier is insufficient to deduct the fees of all the billable items:

When you upload data to the STANDARD storage class, [storage usage fees](#) will be incurred, which can be deducted from the free tier offered by COS. COS provides multiple separately billed storage classes. The free tier of STANDARD cannot be applied to deduction of fees of other storage classes (such as STANDARD_IA).

Objects that need to be accessed or downloaded by other users incur [traffic](#) and [requests](#) fees, which cannot be deducted from the free tier.

If you upload and download data between the COS bucket and CVM instance in the same region, the access is over the private network in the same region and doesn't incur traffic fees (cross-region private network access will incur cross-region traffic fees). If your data is downloaded to the local file system through the console, API, or COS tools, public network downstream traffic fees will be incurred.

COS offers multiple storage classes that generate different types of requests. For example, if STANDARD_IA data is downloaded, STANDARD_IA requests will be generated and billed on a pay-as-you-go basis.

If your COS bucket is used together with CDN, CDN origin-pull traffic fees may also be incurred. For more information, see [Traffic Fees](#).

If you enable global acceleration, global acceleration fees will also be incurred.

Resource usage in excess of the free tier:

For example, if COS offers you free STANDARD storage usage of 50 GB, but your actual usage is 60 GB, the excess of 10 GB will be pay-as-you-go.

Expiration of the free tier:

COS offers new users a free tier of STANDARD storage usage valid for six months. After the free tier expires, storage usage will be pay-as-you-go.

Does the free tier apply to the INTELLIGENT TIERING storage class?

No. The free tier is applicable only to **STANDARD storage usage** but not **other billable items** such as STANDARD_IA/ARCHIVE storage usage, requests, and traffic.

INTELLIGENT TIERING is an independent storage class and incurs INTELLIGENT TIERING storage usage fees, which cannot be deducted from the free tier for the STANDARD storage class. For more information, see [Free Tier](#).

Does COS offer a free tier for CDN usage?

No. COS and CDN are different products. The CDN origin-pull traffic generated by the use of CDN is billed by COS on a pay-as-you-go basis, and the generated CDN traffic is billed by CDN. For more information on the differences between these two types of traffic, see [Traffic Fees](#).

Resource Packs

What is the Resource Package (Prepaid) mode?

COS has two billing modes: Pay-as-You-Go (Postpaid) and Resource Package (Prepaid). Compared to Pay-as-You-Go (Postpaid), the Resource Package (Prepaid) mode offers more benefits. For information on supported types of resource packages, see [Resource Package Overview](#).

How can I purchase a resource package?

You may purchase resource packages following the guide in [Purchasing Resource Packages](Purchasing a Resource Pack!6e6986405f44365d5173872cb56b0214).

Does the specification of the COS traffic package indicate the available quota each month or the total available quota during the validity period?

The specification of the traffic package refers to **the quota available for deduction each month**. Fees of excess usage will be deducted from the account balance, and the quota that is not used up in the current month will not be retained to the next month. If you purchase two traffic packages with the same validity period at the same time, both specifications are counted. The monthly quota is the total quota of the two traffic packages, and fees of excess usage will be deducted from the account balance.

Why are fees still deducted from the account balance or overdue payments incurred even if I have purchased a resource package?

If deduction from the account balance or overdue payments still occur after you purchase a resource package, the possible reasons are as follows:

1. The use of COS involves a variety of billing items, but the resource package you purchase cannot deduct fees of all the billing items.

After your data is uploaded to a COS bucket, [storage usage fees](#) are incurred. You need to purchase a **storage capacity package** for fee deduction. COS provides multiple storage classes, each with the corresponding and specific storage capacity package. For example, a capacity package for standard storage cannot be used to deduct storage fees of low-frequency storage.

If your data in COS buckets can be accessed or downloaded by other users, [traffic fees](#) and [request fees](#) are incurred. You need to purchase a **public network downstream traffic package** and a **data request package** for fee deduction.

If data upload and download are between COS buckets and CVM instances in the same region, the access is deemed intra-regional network access, and no traffic fees are incurred. (Cross-regional network access will incur cross-regional traffic fees.) However, public network downstream traffic fees are incurred when the data is downloaded to local paths through the console, APIs, or COS tools.

COS provides various storage classes, with corresponding types of requests. For example, downloading low-frequency storage data requires low-frequency requests. You must use the request packages corresponding to the specific request types. Currently, only standard request packages and low-frequency request packages are provided. Requests of other storage classes can adopt the pay-as-you-go mode.

If you have configured CDN or EdgeOne for COS buckets, CDN origin-pull traffic fees may be incurred. You need to purchase a **CDN origin-pull traffic package** for fee deduction. For detailed information on the traffic fees incurred when COS is used together with CDN or EdgeOne, see [Traffic Generated by Using COS as CDN Origin](#).

If you have enabled the global acceleration feature, related fees will be incurred. You need to purchase a **global acceleration traffic package** for fee deduction.

2. The region of the resource package you have purchased is not the same as that of your COS bucket. Therefore, the resource package cannot be used for fee deduction.

When you purchase a resource package, you need to select the region of the resource package. Packages with different regions cannot be used concurrently. For example, if you have purchased a standard storage capacity package for regions in the Chinese mainland (excluding the financial cloud region), but your data is stored in a storage bucket in the Singapore region, the fees of this bucket cannot be deducted using the resource package you have purchased.

Currently, COS only provides resource packages for regions in the **Chinese mainland**, the **Hong Kong (China) region**, and **regions outside Chinese mainland**. The incurred fees of COS buckets in other regions (such as financial cloud region) are still billed based on the actual usage. For more details on regions, see [Regions and Access Endpoints](#).

3. Your usage has exceeded the specification of the purchased resource package.

For example, you have purchased a standard storage capacity package of 100 GB, but the actual usage has reached 105 GB. The excess part of 5 GB is charged in the pay-as-you-go mode, which incurs the fees. You can upgrade your resource package to solve this problem.

Are the fees deducted automatically after I purchase a COS resource package, or are there other configuration operations required?

After you purchase a COS resource package, fees will be deducted automatically starting from the effective time you selected. No other configuration operations are needed.

If I have multiple resource packages under my account, will another resource package be automatically used after one expires?

When a resource package expires or runs out, the system will automatically use other available resource packages. No configuration or data migration operations are required. If no other resource packages are available, fees will be deducted from the account balance by default. If the account balance is less than 0 and is not topped up in time, the service will be suspended due to overdue payment.

Can multiple COS resource packages be used concurrently?

COS resource packages can be used concurrently. The specifications of all resource packages are counted, but the validity periods are not. For descriptions of resource packages and related cases, see [Resource Package Overview](#).

Do I need to purchase a CDN origin-pull traffic package if I have already purchased a public network downstream traffic package?

CDN/EdgeOne origin-pull traffic is generated only when you enable CDN and access COS buckets using CDN for domain name acceleration. The public network downstream traffic package does not cover the CDN/EdgeOne origin-pull traffic. If you have enabled CDN acceleration, it is recommended that you purchase a CDN origin-pull traffic package.

What is the difference between renewing and repurchasing COS resource packages?

Renewal: It refers to extending the validity period of the resource package. For example, you have purchased a public network downstream traffic package of 50 GB with a validity period of three months. If you manually renew it for three months, then the validity period of this traffic package will be extended for three months. During the validity period, you will enjoy a traffic quota of 50 GB each month.

Repurchase: It refers to purchasing the packages you want again on the [COS resource package purchase page](#).

Note

If the resource package price changes (for example, the price drops), the resource packages you have purchased cannot be renewed. Instead, you can purchase new packages.

The **free tier** resource packages displayed on the package management page are not renewable. You can purchase new packages.

Why some COS resource packages cannot be renewed or upgraded?

The resource packages under the following circumstances do not support renewal or upgrade. You can purchase new packages on the [COS resource package purchase page](#).

The standard storage capacity packages **provided free of charge** to new users do not support renewal or upgrade. (Such packages include those of 50 GB or 1 TB in all regions as displayed on the resource package management page.)

If the resource package price changes (for example, the price drops), the resource packages you have purchased cannot be renewed or upgraded. You can purchase new packages.

Storage capacity packages support upgrade operations while traffic packages and request packages do not. If you want to upgrade traffic or request packages for a higher specification, you can go to the [COS resource package purchase page](#) to purchase new resource packages. Multiple traffic or request packages can be used concurrently.

Note

For guidance on renewal and upgrade, see [Renewing Resource Packages](#) and [Upgrading Resource Packages](#).

How can I view the usage of resource packages?

You can log in to the **COS console** and choose [Package Manage > Purchased resource packages](#) to view the effective time, expiration time, and usage of purchased resource packages.

Do the COS resource packages support renewal or upgrade?

Yes. See [Renewing Resource Packages](#) and [Upgrading Resource Packages](#) for related operation guides.

What is the difference between renewal and upgrade of COS resource packages?

Resource package renewal refers to extending the validity of a resource package.

Applicable scenarios for resource package renewal: The validity period of the purchased resource package is about to end, requiring an extension.

Case: Assume that you've purchased a public network downstream traffic package of 50 GB with a validity period of three months. If you manually renew the package for another three months, then the validity period of this package will be extended by three months. During the validity period, you will enjoy a traffic quota of 50 GB each month.

Resource package upgrade refers to the increase of the resource package specifications.

Applicable scenarios for resource package upgrade: The purchased resource package no longer satisfies the actual business requirements. Packages with a higher specification are required.

Case: Assume you've already purchased a public network downstream traffic package of 50 GB with a validity period of three months. If you upgrade the resource package specification to 100 GB on the upgrade page, then this package can deduct fees of 100 GB of public network downstream traffic each month, but the validity period remains unchanged.

Can I request a refund on the COS resource packages?

For resource packages that meet the refund criteria, you can apply for a refund. See [Refund on Resource Packages](#) for more information.

Will data be lost after the COS resource package expires? Do I need to migrate the data?

The resource package is used to deduct fees incurred during service usage but not to store data. Once the resource package is used up or expires, **the billing mode changes to the pay-as-you-go (post-paid) mode automatically**, which means that fees are deducted from the account balance. The data will not be lost, and **you do not need to migrate the data**. The data will be terminated only when your account has overdue payment for consecutive 15 days. For details, see [Payment Overdue](#).

What is the relationship between COS resource packages and buckets?

COS bucket provides the storage space for data, while resource package is a prepayment method. Resource packages can deduct fees according to billing items and regions and have no direct relationship with COS buckets. For details, see [Resource Package Overview](#).

For example, you have a standard storage capacity package of 50 GB in a region in the Chinese mainland with a validity period of three months under your account. In that case, all standard storage capacity fees of COS buckets in

regions in the Chinese mainland under your account can be deducted using this resource package.

What resource package should be purchased to deduct the fees incurred if global acceleration is enabled for COS?

The traffic fees incurred by using the global acceleration feature can be deducted by the **global acceleration traffic package**. For the description of global acceleration, see [Overview](#).

Will the expiration of the purchased COS resource package result in service suspension?

Once the resource package expires, the **pay-as-you-go** billing mode will take effect. If your account has overdue payment, the COS service will be suspended 24 hours later. Your data can still be retained for another 15 days. If the balance is not topped up during this period, your data will be terminated. For more information, see [Payment Overdue](#).

Can I transfer the COS resource packages to another Tencent Cloud account for use?

COS resource packages cannot be transferred to other accounts but only be used under the account that has purchased the packages.

Can COS resource packages deduct request fees?

Different billing items of COS correspond to different types of resource packages. Request packages are required for deducting request fees. Currently, COS provides **standard storage request packages** and **low-frequency storage request packages**. You can purchase the request packages on the [COS resource package purchase page](#).

Do I need to purchase traffic packages for using COS?

You can decide whether to purchase traffic packages based on the actual scenario. COS has multiple billing items, and the free tier resource packages can only deduct the standard storage capacity fees. During the use, other fees such as traffic fees and request fees may be incurred. It is recommended to purchase the **public network downstream traffic packages** and **request packages** before you use COS to avoid deduction from your account balance. You can go to the [COS resource package purchase page](#) to purchase the packages.

Can purchased COS resource packages deduct overdue payment?

If the purchased resource package **takes effect immediately**, the fees incurred on the same day can be deducted using the package. The fees incurred before the purchase cannot be deducted. Therefore, service packages cannot deduct overdue payment.

How are COS resource packs applied to deduction?

When bills are settled within the validity period of the COS resource package, the following rules apply:

Resource package type: Free tier resource packages are used for deduction before the purchased resource packages. If the total usage exceeds the specifications of both free tier and purchased resource packages, the excess part is billed in pay-as-you-go mode.

Expiration time of resource package: Resource packages that expire first will be used first. Note: **The expiration time here refers to the time specified in the validity period of the resource package instead of that of the current cycle.**

Resource package usage: Resource packages with more remaining quotas will be used first.

Purchase time of resource package: Resource packages that are purchased first will be used first.

For further description of the rules, see the [Types and Rules](#) section in Resource Package Overview.

What are the unit conversion rules of COS resource packages?

COS provides storage capacity packages, traffic packages, and request packages.

The billing unit of the storage capacity billing item and storage capacity packages is GB, and the capacity unit conversion ratio is 1:1,024. For example, 1 TB of standard storage capacity equals 1,024 GB. A standard storage capacity package of 1 TB can deduct fees of 1,024 GB of standard storage capacity.

The billing unit of the traffic billing item and traffic packages is GB, and the traffic unit conversion ratio is 1:1,024. For example, 1 TB of public network downstream traffic equals 1,024 GB. A public network downstream traffic package of 1 TB can deduct fees of 1,024 GB of public network downstream traffic. This conversion rule will take effect from April 1, 2025.

The billing unit of the request billing item and request packages is 10,000 requests. For example, A standard storage request package with a quota of 100,000 requests can deduct fees of 100,000 standard storage requests.

What are the expiration, isolation, and termination of a COS resource package?

1. Applicable scope:

The expiration, isolation, and termination policies of COS resource packages only apply to purchased packages (those in the list on the [Purchased resource packages](#) page) but not to the packs provided free of charge (those in the list on the [Free tier resource package](#) page). You can view the type of your resource packages on the [Package Manage](#) page in the COS console.

2. Definitions

Expiration: It refers to the end of the validity period of your resource packages.

You can view the package validity period on the [Package Manage](#) page in the COS console. "Effective" indicates the start time of the validity period, and "Expires on" indicates the end time of the validity period.

You can view the expiration date of the packages on the [Renewal Management](#) page in Billing Center.

Isolation: It refers to the condition in which your resource package will be isolated one day after it expires.

You can view the isolated packages in the list of expired packages on the [Package Manage](#) page in the COS console.

Isolated packages are not displayed on the [Renewal Management](#) page in Billing Center and cannot be renewed. If such resource packages are still needed, you can [purchase again](#).

Termination: Your resource package will be isolated one day after it expires and will be terminated after it is isolated for 365 days.

Terminated packages are not displayed on the [Package Manage](#) page in the COS console.

Terminated packages are not displayed on the [Renewal Management](#) page in Billing Center.

Notification of changes in expiration, isolation, and termination policies for COS resource packages

Details:

Applicable scope

The changes in the isolation and termination policies apply only to purchased resource packages (those in the list on the [Purchased resource packages](#) page) instead of free resource packages (those in the list on the [Free tier resource package](#) page). You can view the type of your resource packages on the [Package Manage](#) page in the COS console.

Policy Description

| Policy | Before Change | After Change |
|-------------------------------------|--|--|
| Expiration of COS resource package | A notification is sent both before and after the resource package expires. | A notification is sent both before and after the resource package expires, but the notification content is updated. |
| Isolation of COS resource package | No isolation policy. | The resource package is isolated one day after it expires. Isolated resource packages are displayed on the Package Manage page in the COS console but not on the Renewal Management page in the Billing Center. |
| Termination of COS resource package | No termination policy. | The resource package is isolated one day after it expires and terminated after it is isolated for 365 days. Terminated resource packages are not displayed on the Package Manage page in the COS console or the Renewal Management page in the Billing Center. |

Case

If you purchased a standard storage package of 10 GB on January 1, 2023 (UTC+8) with a validity period of one month, this package will expire on January 31, 2023 (UTC+8). You will receive a notification on the expiration both before and after the package expires.

Before the change: After the resource package expires, you can find it on the Package Manage page in the COS console and the Renewal Management page in Billing Center. The package still exists and is not terminated.

After the change: The resource package is isolated one day after it expires. This resource package is displayed on the Package Manage page in the COS console but not on the Renewal Management page in Billing Center. After the resource package is isolated for 365 days, it is terminated and no longer displayed on the Package Manage page in the COS console or the Renewal Management page in Billing Center.

Notification

What types of notifications does COS send?

COS sends notifications for new feature launch, product change, expiration, repossession, and alarms.

What are subscribed and non-subscribed messages?

Subscribed messages: You can customize their recipients and notification methods in Message Center. They are used in default scenarios.

Non-subscribed messages: You cannot modify their recipients and notification methods but can only receive them passively. They are used in special scenarios.

How do I configure the root account and sub-accounts to receive COS notifications?

You can customize the recipients of COS notifications by setting **Message Recipient** in [Subscription Management](#) in Message Center.

Traffic

What are the differences among COS public network downstream traffic, CDN origin-pull traffic, and global acceleration traffic?

Public network downstream traffic: The traffic generated by data transfer from COS to the client over the internet, such as the traffic generated when you directly access a resource at a COS domain name in a browser.

CDN origin-pull traffic: The traffic generated by data transfer from COS to a CDN edge node when you access a CDN acceleration domain after enabling CDN acceleration.

Global acceleration traffic: The traffic generated by data transfer with acceleration domains after global acceleration is enabled.

How is the public network downstream traffic in COS generated and billed?

Public network downstream traffic is traffic generated by data transfer from COS to the client over the Internet. Traffic generated by downloading an object directly through an **object link** or by browsing an object through a **static website access node** is public network downstream traffic. For more information on billing for such traffic, see [Billable Items](#) and [Product Pricing](#).

Will I be charged for public network downstream traffic generated by downloading files through the COS console, tools, API, or SDK?

The traffic (private or public network traffic) generated by accessing COS is subject to the use case, and only access to COS from a Tencent Cloud product in the same region will be over the private network by default, with no public network downstream traffic fees incurred. For more information on how to identify private network access, see [Overview > Private Network Access](#).

What is public network traffic in COS?

Public network downstream traffic is the traffic generated by data transfer from COS to the client over the internet. Downloading a file stored in COS in the COS console, accessing or downloading an object through a tool, object address, or custom domain name, and previewing an object in a browser will generate public network downstream traffic. For more information, see [Overview](#).

Will accessing COS over the private network incur fees?

Accessing COS over the private network will incur **storage usage fees** and **request fees** but not **traffic fees**. For more information, see [Billable Items](#).

How will COS be billed after it is connected to CDN?

After COS is connected to CDN, the fees incurred by COS and CDN will be billed separately. COS fees include storage usage fees, request fees, and CDN origin-pull traffic fees. CDN fees include CDN traffic fees.

Why is there public network downstream traffic after I enabled CDN acceleration?

It may be because you are still using the COS origin domain name formatted as `<BucketName-APPID>.cos.<region>.myqcloud.com` to access COS files after you enabled CDN acceleration. We recommend that you use the CDN acceleration domain name instead, which only generates CDN origin-pull traffic.

What is CDN origin-pull traffic in COS? How is it generated?

CDN origin-pull is to pull data from COS to the cache node by CDN when a file not cached on a CDN edge node is accessed at a CDN domain name.

Note:

CDN origin-pull will incur origin-pull traffic fees. For more information, see [Traffic Fees](#).

How is the CDN origin-pull traffic billed in COS?

CDN origin-pull traffic is the traffic generated by data transfer from COS to a Tencent Cloud CDN edge node. After CDN acceleration is enabled, traffic generated by browsing or downloading COS data on the client at a **CDN acceleration domain name** is CDN origin-pull traffic. For billing details, see [Billing Overview](#) and [Pricing | Cloud Object Storage](#).

What are the differences between CDN origin-pull traffic and CDN traffic?

CDN origin-pull traffic is a billable item in COS. It is the origin-pull traffic generated by data transfer from COS to the CDN edge node when COS is used as the CDN origin.

CDN traffic is a billable item in CDN. It is the traffic generated by data transfer from a CDN edge node to the client.

Will fees be charged for traffic and requests generated by data transfer between COS and CVM?

When data transfer happens between COS and CVM from different regions, there are charges for both traffic and requests. When data is transferred in the same region, there are charges for requests, but not for traffic (free for data transfer over private network). For more information, see [How to determine access over private network](#).

Will there be traffic fees when I upload a file to a COS bucket?

No. The upstream traffic generated by file uploads is free of charge.

Will there be traffic fees when I access Tencent Cloud products in the same region?

Tencent Cloud products within the same region access each other over the private network by default, with no traffic fees incurred. For more information, see [How to determine an access over private network](#).

What is traffic anti-fraud in COS?

1. Feature Overview

COS supports the Hotlink Protection feature. The system intelligently detects abnormal traffic, automatically blocks abnormal buckets, and sends notifications. To avoid a high bill due to abnormal traffic before auto-blocking, refer to [Hotlink Protection Guide](#) to implement protective measures.

Intelligent detection is described as follows:

The system makes comprehensive judgment based on recent traffic changes in the buckets. Not all traffic increases are considered abnormal.

Abnormal traffic occurs before intelligent detection starts, so it cannot be detected in advance.

If a bucket has abnormal traffic before it is blocked, the traffic will be reflected in the bill, and the fees will be borne by the user.

Upon receiving a notification, whether blocking is caused by abnormal traffic or anticipated traffic increases, please [contact us](#) for unblocking.

Note:

If the bucket access permission is set to Public Read (Public Read/Private Write or Public Read/Write), the hotlink risk is extremely high. It is recommended to take measures to avoid unexpected fees. Refer to [Hotlink Protection Guide](#).

When a bucket is created, COS informs you of the billing risks and receives your confirmation. If fees arise from the Public Read permission for the bucket, no configuration of hotlink protection, etc., they will be borne by the user.

2. Scope of Application

The traffic anti-hotlinking feature is a system policy that automatically carries out intelligent detection for each user.

3. Allowlist

If you do not need the traffic anti-hotlinking feature, after critically assessing the related risks and confirming that the risks are bearable, provide the bucket name and contact business personnel or [contact us](#) to apply for a Traffic anti-hotlinking allowlist. After being added to the allowlist, your bucket will no longer undergo anti-hotlinking detection.

4. Notification Logic

If the system detects abnormal total traffic growth of a bucket for a day, to reduce the high traffic costs caused by unexpected public network downstream traffic, COS will block the bucket and send a notification.

Note:

1. Total traffic of the bucket for the day: total traffic generated in the duration from 00:00 of the day to the detection time.

2. Traffic in traffic anti-hotlinking: billed items in [public network downstream traffic](#).

5. Measures

After the bucket is blocked due to hotlinking, take the following measures to ensure normal access to the bucket:

Self-service unblocking is not supported. After receiving the <Account Blocked for COS Data Hotlinking>, [contact us](#) as soon as possible for unblocking, so as to ensure normal access to the bucket.

You can log in to the [COS console](#) to view the bucket's traffic monitoring and configure the bucket's permissions and alarm policies as needed.

Bill

How do I view my bill?

You can view the fees incurred by the use of COS under your account in the [Billing Center](#) in the console. For more information, see [Viewing and Downloading Bill](#). You can view bucket-level bill details by downloading the usage details in the [Billing Center](#).

How do I download a bill?

Log in to the Tencent Cloud console, select **Billing Center** > [Bill Download Center](#), and download the target bill packages, PDF bills (L0), bill summary (L1), bills by instance (L2), and bill details (L3). For more information, see [Bill Download Center](#).

What are billing by bucket and cost allocation by tag?

Billing by bucket: It refers to using the bucket name and `APPID` together as the **resource ID** and using the bucket name as the **instance name**, i.e., generating bills by bucket. You can view the fees and usage of billable items by bucket.

Cost allocation by tag: You can specify **cost allocation tags** to differentiate resources by category.

How do I set billing by bucket and cost allocation by tag?

COS supports billing by bucket and cost allocation by tag.

Billing by bucket: Contact your sales rep to add your account to the allowlist of the billing by bucket feature. This feature will take effect for new bills the next day after approval, while historical bills will remain unchanged.

Cost allocation by tag: Tag your buckets, set cost allocation tags, and contact your sales rep to add your account to the allowlist of the billing by bucket feature.

Note:

After your account is added to the allowlist, billing by bucket will take effect only for new bills, while existing bills will remain unchanged. In addition, you cannot go back to the previous billing option.

How do I determine whether bills are generated by bucket?

You can determine this in two ways:

Option 1: Select **Billing Center** > **Billing Details** and view **Bill by Instance** and **Bill Details**. If the **Resource Alias/ID** is a bucket name and APPID, billing by bucket has been enabled.

Option 2: Select **Billing Center** > **Bill Download Center** and download bill details (L3). If the **Resource ID** is a bucket name and APPID and **Instance Name** is a bucket name, billing by bucket has been enabled.

Below are the effects after billing by bucket is enabled:

(1) Bill by instance

(2) Bill details

(3) Bill details (L3)

Note: After the billing by bucket feature of COS is fully launched, on the **Cost Allocation - Month** sheet in the usage details, the bucket cost allocation information won't be displayed, while other information stays unchanged. For specific bucket fees, download the L2 or L3 bill.

How is the billing by bucket feature of COS launched?

The billing by bucket feature of COS was made available through an allowlist starting from October 28, 2022 as follows:

If you had applied to add your UIN to the allowlist of the billing by bucket feature, the **Resource ID** and **Instance Name** columns in the bill would change, where the **Resource ID** is a bucket name + `APPID` and **Instance Name** is a bucket name.

If you hadn't applied to add your UIN to the allowlist of the billing by bucket feature, billing would stay unchanged.

Starting from February 1, 2023, the feature was **made available to users in batches by UIN** as follows:

| Launch Date | Description | Bill Description |
|-------------------|--|---|
| October 28, 2022 | Feature launch - through allowlist. | (1) Billing stayed unchanged for users not in the allowlist. (2) Billing by bucket was enabled for users in the allowlist. |
| February 1, 2023 | Before full launch - allowed account verification. | (1) Billing stayed unchanged for users not in the allowlist. (2) Billing by bucket was enabled for users in the allowlist. |
| February 6, 2023 | Batch 1 launch - for UINs ending in 9 | (1) Billing by bucket was enabled for users in the allowlist and with UINs ending in 9. (2) Billing stayed unchanged for other users. |
| February 20, 2023 | Batch 2 launch - for UINs ending in 2 | (1) Billing by bucket was enabled for users in the allowlist and with UINs ending in 9 or 2. (2) Billing stayed unchanged for other users. |
| March 7, 2023 | Batch 3 launch - for UINs ending in 3 | (1) Billing by bucket will be enabled for users in the allowlist and with UINs ending in 9, 2, or 3. (2) Billing will stay unchanged for other users. |
| March 14, 2023 | Batch 4 launch - for UINs ending in 4 | (1) Billing by bucket will be enabled for users in the allowlist and with UINs ending in 9, 2, 3, or 4. (2) Billing will stay unchanged for other users. |
| March 21, 2023 | Batch 5 launch - for UINs ending in 5 | (1) Billing by bucket will be enabled for users in the allowlist and with UINs ending in 9, 2, 3, 4, or 5. |

| | | |
|--|--|--|
| | | (2) Billing will stay unchanged for other users. |
|--|--|--|

How do I view the billing statistical period?

Log in to the Tencent Cloud console, select **Billing Center** > [Bill Overview](#), and view the billing statistical period of your account.

What are billing by deduction cycle and billing by billing cycle?

Billing by deduction cycle: The system generates a bill per calendar month based on the **resource fees deduction time**.

Billing by billing cycle: The system generates a bill per calendar month based on the **actual resource usage time**.

What is the relationship between the billing mode and billing statistical period?

COS adopts the pay-as-you-go (postpaid) billing mode.

Pay-as-you-go (postpaid) resources

Daily settled resources: Fees incurred from 00:00 to 23:59 on January 31 will be deducted on February 1. The record will be posted to the bill for February by deduction cycle and to the bill for January by billing cycle.

Monthly settled resources: Fees incurred from 00:00 on January 1 to 23:59 on January 31 will be deducted on February 1. The record will be posted to the bill for February by deduction cycle and to the bill for January by billing cycle.

For more information on billing cycles, see [Bill Management](#).

Why did the amount of the bill (by deduction cycle) of the first month “increase” after the upgrade from monthly to daily settlement?

Starting from July 1, 2022, the settlement cycle of COS storage usage, request, and data retrieval fees was upgraded from monthly to daily to help you manage fees in a more refined manner. The upgrade was implemented on user accounts in batches and went through a two-month beta test. The release dates of different bill statistical periods are as listed below. For more information, see [Daily Billing for COS Storage Usage, Request, and Data Retrieval](#) and [Bill Management](#).

| Release Date | Release Note | Bill Description |
|--------------|--|--|
| July 1, 2022 | The first release for the first batch of accounts in beta test | (1) Resources were settled monthly before July 1, 2022. Fees incurred from 00:00 on June 1 to 23:59 on June 30 were deducted on July 1. The record was posted to the bill for July by deduction cycle and to the bill for June by billing cycle. |

| | | |
|-------------------|--|---|
| | | (2) Resources were settled daily after July 1, 2022. Fees incurred from 00:00 to 23:59 on July 1 were deducted on July 2. The record was posted to the bill for July 2 by deduction cycle and to the bill for July 1 by billing cycle. |
| August 1, 2022 | The second release for the second batch of accounts in beta test | <p>(1) Resources were settled monthly before August 1, 2022. Fees incurred from 00:00 on July 1 to 23:59 on July 31 were deducted on August 1. The record was posted to the bill for August by deduction cycle and to the bill for July by billing cycle.</p> <p>(2) Resources were settled daily after August 1, 2022. Fees incurred from 00:00 to 23:59 on August 1 were deducted on August 2. The record was posted to the bill for August 2 by deduction cycle and to the bill for August 1 by billing cycle.</p> |
| September 1, 2022 | The third release for all accounts | <p>(1) Resources were settled monthly before September 1, 2022. Fees incurred from 00:00 on August 1 to 23:59 on August 31 were deducted on September 1. The record was posted to the bill for September by deduction cycle and to the bill for August by billing cycle.</p> <p>(2) Resources were settled daily after September 1, 2022. Fees incurred from 00:00 to 23:59 on September 1 were deducted on September 2. The record was posted to the bill for September 2 by deduction cycle and to the bill for September 1 by billing cycle.</p> |

Therefore, if the monthly usage remained the same, after monthly settlement was upgraded to daily settlement, the bill amount varied by billing statistical period and billable items as follows:

Billing by billing cycle: The monthly fees of storage usage, requests, and data retrievals remain basically unchanged after the upgrade.

Billing by deduction cycle: The monthly fees of storage usage increase after the upgrade, while the monthly fees of requests and data retrievals remain basically unchanged after the upgrade.

In the first month after the upgrade from monthly settlement to daily settlement, the storage fees increased. This was because two bills were generated for billing by deduction cycle. The first bill was the monthly bill for the last month, while the second bill was the daily bill for the current month. Therefore, the bill amount seemed to have increased, but no additional fees were deducted in fact, which was normal under the settlement and billing logic.

In the second month after the upgrade from monthly settlement to daily settlement, your bills were settled daily, and the bill amount "decreased" compared with that in the first month, which was also normal under the settlement and billing logic.

For example:

Your account is billed by deduction cycle and was upgraded from monthly settlement to daily settlement on September 1, 2022, and you downloaded the **bill details (L3)** of COS in the [Billing Center](#) on September 30.

Taking the **COS STANDARD storage usage** billable item as an example, the bills consisted of the monthly bill for August and daily bill for September as detailed below:

Pay-as-you-go monthly settlement: The monthly bill was generated on September 1 for resource usage fees incurred in the entire August (from 00:00 on August 1 to 23:59 on August 31).

Pay-as-you-go daily settlement: After the upgrade from monthly settlement to daily settlement on September 1, a daily bill was generated on each day starting from September 2 for daily resource usage fees incurred in the entire September. The bill generated on September 2 was for fees incurred from 00:00 to 23:59 on September 1, that on September 3 was for fees incurred from 00:00 to 23:59 on September 2, and so on.

As there was too much bill data, only the daily bills for September 2–3 are listed below:

Therefore, if your monthly fees increased in the first month after the upgrade from monthly settlement to daily settlement and your bill was generated by deduction cycle, the increase was normal under the settlement and billing logic, and no additional fees were deducted in fact. If you have any questions, [contact us](#) for assistance.

Why are there daily and monthly settlement modes in the transaction details?

Before the upgrade, COS' storage, requests, STANDARD_IA data retrievals, and ARCHIVE data retrievals are settled monthly, and their monthly settlement details are displayed in the transaction details, while DEEP ARCHIVE data retrievals, traffic, and management features are still settled daily, and their daily settlement details are still displayed in the transaction details. After the upgrade, all COS billable items are settled daily. For more information, see [Billing Overview](#).

In addition, you can view the L3 bill to check whether the billable items are settled daily or monthly as instructed in [Billing](#).

Fee Deduction

Will I be charged immediately after activating COS?

Activating COS is free of charge, and fees will be incurred only after you use it. For billing details, see [Billing Overview](#).

If you upload multiple files in a total size of 1 GB to the STANDARD storage class after activating COS and then view and download them, STANDARD storage, write requests, public/private network upstream traffic, and public network downstream traffic (of about 1 GB) will be generated, and the bill generated on the next day will include the STANDARD storage, STANDARD write request, and public network downstream traffic fees. As the public/private network upstream traffic is free of charge, it will not be included in the bill.

Will file uploads to COS incur fees?

When you upload a file to COS, the generated **traffic is free of charge**, but the generated **write request** will incur fees, and the **storage usage fees** will be calculated based on the file size. If you access or download the object over the public network, public network downstream traffic fees will be incurred. For billing details, see [Billing Overview](#) and [Pricing | Cloud Object Storage](#).

Why are fees still deducted after my data in COS is deleted?

If you no longer use the COS service, you need to delete all the buckets under your account. Double check whether **all the buckets have been deleted**. If fees are still incurred after all the buckets are deleted, they may be **the daily fees incurred on the previous day**. COS storage usage fees and request fees are billed daily, that is, a bill generated on the current day is for usage on the previous day. You can go to the [transaction details](#) page and click **Details** on the right of a billable item to view fee deduction details.

The details about the change from monthly settlement to daily settlement are as follows: Starting from July 1, 2022, the settlement cycle of COS storage usage, request, and data retrieval fees was upgraded from monthly to daily to help you manage fees in a more refined manner. The upgrade was implemented on user accounts in batches and went through a two-month beta test. The release dates of different bill statistical periods are as listed below. For more information, see [Daily Billing for COS Storage Usage, Request, and Data Retrieval](#) and [Bills](#).

How do I deactivate the COS service and stop being charged?

You can deactivate COS or stop its billing as follows:

1. If you decide to stop using COS, you can avoid any further billing by permanently deleting all of your COS data (including incomplete multipart uploads and object versions) as instructed in [Payment Overdue](#). There is no need to remove your account, and if you use other Tencent Cloud products, avoid doing so as it will affect your other services.
2. If you don't use COS for more than one month, you can set lifecycle rules to transition data in STANDARD storage class in the bucket to a colder class such as STANDARD_IA, ARCHIVE, or DEEP ARCHIVE to reduce storage fees. For more information, see [Setting Lifecycle](#). The transition will generate read requests in the original storage class and write requests in the target storage class, so transition by lifecycle will incur read/write [request fees](#).

Notes

Data, once deleted from the bucket, cannot be recovered, so make backups accordingly.

If versioning is enabled for the bucket, disable it before deleting data.

Check your billing cycle to avoid overdue payments. If all your billable items are settled daily, the bill for the day of data deletion will be generated on the next day. After the data is completely cleared, the system will stop billing. For more information, see [Billing Cycle](#).

If your account has overdue payment (i.e., your account balance is below 0), COS services will be suspended after 24 hours, regardless of whether your resource pack is within the validity period.

If your account has overdue payment and COS services are suspended, the free tier for which your account is eligible won't be available.

If data in your bucket is blocked for the second time due to non-compliance, it cannot be deleted. [Contact us](#) if you have any questions.

How will I be charged when storing data in STANDARD_IA for less than 30 days?

The minimum data storage duration in STANDARD_IA storage class is 30 days.

Below are the specific rules:

An object stored less than 30 days is billed by 30 days as described in [Storage Usage Fees](#).

An object stored more than 30 days is billed by the actual storage duration.

For more information, see "STANDARD_IA storage usage fees" and "Billing Mode and Calculation Formula" in [Storage Usage Fees](#).

How will I be charged when storing data in ARCHIVE for less than 90 days?

The minimum data storage duration in ARCHIVE storage class is 90 days.

Below are the specific rules:

An object stored less than 90 days is billed by 90 days as described in [Storage Usage Fees](#).

An object stored more than 90 days is billed by the actual storage duration.

For more information, see "ARCHIVE storage usage fees" and "Billing Mode and Calculation Formula" in [Storage Usage Fees](#).

How will I be charged when storing data in DEEP ARCHIVE for less than 180 days?

The minimum data storage duration in DEEP ARCHIVE storage class is 180 days.

Below are the specific rules:

An object stored less than 180 days is billed by 180 days as described in [Storage Usage Fees](#).

An object stored more than 180 days is billed by the actual storage duration.

For more information, see "DEEP ARCHIVE storage usage fees" and "Billing Mode and Calculation Formula" in [Storage Usage Fees](#).

What are data retrieval fees in COS?

Data retrieval fees are fees incurred by reading or downloading **MAZ_STANDARD_IA/STANDARD_IA** data or restoring **ARCHIVE** or **DEEP ARCHIVE** data to STANDARD. They are calculated based on the amount of retrieved data. The higher the amount, the higher the fees. If you don't have special storage needs, you can directly use STANDARD, which doesn't involve data retrieval fees.

What fees will be incurred by copying MAZ_STANDARD_IA/STANDARD_IA data?

Copying MAZ_STANDARD_IA/STANDARD_IA data will incur request fees and data retrieval fees and may also incur cross-region replication traffic fees if the destination and source files are in different regions.

For the calculation details of such fees, see [Request Fees](#), [Data Retrieval Fees](#), and [Traffic Fees](#).

Will I be charged for a copy generated by restoring ARCHIVE or DEEP ARCHIVE data in COS?

A copy generated by restoring ARCHIVE or DEEP ARCHIVE data is in the **STANDARD** storage class and will incur STANDARD storage usage fees.

How will I be charged for less than 10,000 read/write requests?

The read/write requests are priced by storage class and billed based on the actual quantity, with a minimum billable quantity of 10,000. For more information, see [Pricing | Cloud Object Storage](#) and [Request Fees](#).

Why are the COS read/write request fees zero?

If the number of requests does not reach the minimum value for fee deduction, the request fees will be zero.

Case study: If 23 STANDARD read requests for your data stored in the STANDARD storage class in a bucket in Beijing region were made in December 2021, the unit price of STANDARD read/write requests was 0.002 USD/10,000 requests, and your account wasn't entitled to any discount, then the STANDARD read request fees were $0.0023 * 0.002 = 0.0000046$ USD. As fee deduction is accurate down to two decimal places, your request fees for the month were 0 USD.

Cause: As bills support eight decimal places at most, while fee deduction is accurate down to two decimal places, the system will automatically adjust the accuracy difference. For more information, see [Bills](#).

Why Are There STANDARD Storage Request Fees in My Bill When MAZ_STANDARD in COS Is Used?

While the storage pricing for the STANDARD and MAZ_STANDARD storage in the same region differs, the pricing for read/write requests remains the same. If you use MAZ_STANDARD and some bucket-related requests are generated, the system calculates the usage based on STANDARD read/write requests. As a result, STANDARD storage read/write request fees will be incurred in your bill.

Overdue Payment and Service Suspension

Can I still access and download files in the COS console if my COS service is suspended due to overdue payment?

After the COS service is suspended due to overdue payment, you cannot read or write data from or to COS, but you can top up your account. For more information, see [Payment Overdue](#).

Why can't I use the COS service after topping up my account to a positive balance?

Your account will be automatically unblocked in 10 minutes after it is topped up. If your account is still displayed as blocked, the browser may have cached the historical page. In this case, we recommend you **refresh the webpage** or **clear the browser cache** first.

Notice About the Change in Service Suspension Policy for Overdue COS Pay-As-You-Go (Postpaid) Payment

For details, see [Notice About the Change in Service Suspension Policy for Overdue COS Pay-As-You-Go \(Postpaid\) Payment](#).

Data Processing

How is the PowerPoint-to-image conversion in COS' document processing service billed? Will fees be incurred if I access an object that I have accessed before?

The document processing service is offered by CI, so its fees are also charged by CI. Fees will be incurred for every request; that is, fees will be deducted every time you refresh the URL. For more information, see [Document Processing Fees](#).

Others

How will I be charged for migrating data from another cloud to COS?

When you migrate data from another cloud to Tencent Cloud COS, outbound traffic fees will be charged by your source cloud storage vendor. The write traffic generated by migration to Tencent Cloud is free of charge, but storage usage and request fees will be incurred. For COS billing details, see [Billing Overview](#).

What fees will be incurred by accessing COS through a URL?

[Public network downstream traffic fees](#) and [request fees](#) may be incurred. If you enable CDN and access data through a CDN domain name, CDN traffic fees and [CDN origin-pull traffic fees](#) will also be incurred.

Do IOPS, latency, and throughput of COS vary by price?

No. For more information, see [Specifications and Limits](#).
traffic anti-hotlinking allowlist

Bucket Configuration Lifecycle

Last updated : 2024-04-24 16:40:57

How can I modify the storage class for an object?

You can modify the storage class in either of the following ways:

1. Use the console to modify the storage class of one or more objects. For detailed directions, please see [Modifying Storage Class](#).
2. [Set a lifecycle rule](#) to transition objects in a specified bucket or prefixed with a specified value to another storage class.

What should I do if an object I deleted still exists?

This may be caused because the object contains special characters. You can [set a lifecycle rule](#) to delete it.

Can I set an expiration time for my objects in COS?

You can set a lifecycle rule to specify the number of days to transition objects to ARCHIVE or delete them after they are uploaded. For detailed directions, please see [Setting Lifecycle](#). You can also use an SDK to set a lifecycle (see [Android SDK - Lifecycle](#) or [Java SDK - Lifecycle](#)).

Can objects be automatically cleared regularly?

You can [set a lifecycle rule](#) to automatically transition objects to another storage class or delete them.

After enabling versioning, how can I manage historical versions?

You can [set a lifecycle rule](#) and enable **Managing historical versions** to transition objects or delete historical object versions.

How many lifecycle rules can I configure?

You can configure up to 1,000 lifecycle rules for a bucket. For more information, please see [Lifecycle Overview](#).

Note:

It is recommended that you do not configure conflicting lifecycle rules for the same set of objects in COS, because this may result in different fees.

When will a lifecycle rule configured take effect?

Rules configured at any time will always start executing at 00:00 the next day. As objects are executed asynchronously, the execution of qualified objects uploaded after the rule is configured will be completed before 24:00 the next day in most cases.

For example, you configured a lifecycle rule at 15:00 on the 1st day of the month to delete files one day or longer after they are modified. Then, at 00:00 on the 2nd day, the lifecycle task scans for files that were modified over one day ago and deletes them. Files uploaded on the 1st day will not be deleted at 00:00 on the 2nd day, as the time elapsed since their modification is less than one day. Instead, these files will be deleted at 00:00 on the 3rd day.

What does “deleting incomplete multipart” mean in the lifecycle configuration?

In lifecycle configuration, **Deleting incomplete multipart** means to delete incomplete multipart uploads generated by multipart upload. Parts that are uploaded but not merged become incomplete multipart uploads, which take up storage but cannot be downloaded. You can check **Deleting incomplete multipart** as needed.

Batch Operations

Last updated : 2024-12-26 15:03:09

How can I restore archived objects to STANDARD in batches?

You can do so as follows:

1. Enable the [Inventory](#) feature. Then, generate an inventory file for objects that need to be restored to STANDARD, and wait for the inventory file to be generated.
2. Create a task that restores archived objects in batches. When configuring the task, select the inventory file and set the copies' effective period (e.g., 7 days). For detailed directions, please see [Batch Operation](#).
3. As the data volume is large, the restoration may take a long time. You can wait for 48 hours after the task is created for the restoration to complete. After this, you can generate and download the inventory for filtering. You need to remove STANDARD objects in the inventory file and keep only the archived objects. Then, upload the modified inventory file to COS.
4. Create a batch replication task. When configuring the task, select the newly uploaded inventory file, set the storage class to STANDARD, and wait for the task to complete.

Does COS have the batch compression feature?

Currently, COS does not support batch compression. You can use SCF to add a file decompression rule for a bucket to automatically decompress files to a specified bucket and path. For more information, please see [File Decompression](#).

How can I obtain the inventory file for the batch operation task?

You can obtain the inventory file in either of the following ways:

Use [COS's Inventory feature](#) to generate an inventory file. Once the inventory file is generated, you can go to the bucket to pull the `manifest.json` file.

Save files that need to be processed in a local `CSV` file and then upload it to COS. Fields required are shown below. For more information, please see [Inventory Overview](#).

```
Bucket,Key,VersionId
examplebucket-1250000000,testFile.txt,testVersionId
```

Why is my data not restored after the batch archived file restoration task is completed?

After the restoration request is sent, the backend will restore files to STANDARD in sequence according to the restoration mode. However, there is a time difference. The completion message displayed in the frontend only indicates that all restoration requests are sent, but not the completion of the restoration task. You can wait for a while and then log in to the console to view the status.

Origin-Pull

Last updated : 2024-12-26 15:05:37

What is COS origin-pull?

When the data you want to access is not stored in COS, you can use COS's origin-pull feature to pull data from a specified origin server (e.g., a local IDC, or the origin server/bucket of other cloud vendors).

Origin-pull is mainly used for hot data migration, redirection for specified requests, and other scenarios. You can configure it as needed. For detailed directions, please see [Setting Origin-Pull](#).

After an origin-pull rule is set, if the requested object does not exist in the bucket, the object can be found from the configured origin server address using the origin-pull rule and returned to the user. Likewise, when you need to redirect specific requests, the origin-pull rule can be used for COS to access data in the origin server.

How can the client know whether a COS request pulls from an origin server?

If the origin-pull is asynchronous, after origin-pull is configured, 302 will be returned for the first COS request, and the client's second request will be forwarded to the origin server. If the origin-pull is synchronous, COS will pull data from the origin server in real time and return it to the client, and also dump the data to the server.

How will I be notified after the offline origin-pull upload succeeds?

SLA of the offline origin-pull module is not always successful. If you want to know whether the origin-pull upload is successful, you can go to the SCF console to set a callback that is triggered by the offline origin-pull. For more information about SCF's COS triggers, please see [COS Trigger](#).

What is an origin-pull address for?

An origin-pull address is usually an IP or a domain name. It specifies where the data you need to pull is stored. When COS does not have the resource you want to access, you can use the origin-pull address to pull the resource in real time.

After origin-pull is configured, if COS does not have the resource/path corresponding to the origin-pull address, will COS upload the resource and create a path after the user's initial access?

Yes. COS will pull the resource automatically and create a path.

Bucket Tags

Last updated : 2024-12-26 15:06:41

How many tags can be added to a bucket?

Up to 50 bucket tag keys can be added for a single bucket, and each tag key can have up to 1,000 tag values. Besides, each root account can add up to 1,000 tag keys. For more information about bucket tag restrictions, please see [Bucket Tag Overview](#).

How can I use bucket tags for cost allocation?

After you [add a bucket tag](#) for your bucket, tags will be recorded in your monthly [usage bill](#). You can download the bill and then create a pivot table (see [Cost Allocation Tags](#)) to analyze the tag-specific usage for your bucket.

How can I use bucket tags to manage access permissions?

You can do so as follows:

1. Contact the root account to obtain permissions to create buckets (`PutBucket`) and operate resources that have a specified tag.
2. Create a bucket and add tags for it.
3. Operate COS objects with the authorized APIs.

Hotlink Protection

Last updated : 2024-12-26 16:22:02

How can I prevent my COS files from hotlinking?

1. If your files are accessed through a browser, you can use the hotlink protection feature to configure an allowlist or blocklist. For detailed directions, please see [Setting Hotlink Protection](#).
2. If your objects are accessed directly via URLs anonymously, you can [add a bucket policy](#) to set the IP allowlist/blocklist. For more information, please see [Access Policy Language Overview](#).
3. For signed requests, configuring a blocklist or allowlist is not supported.

What do I do if the hotlink protection configuration does not take effect when I use a CDN acceleration domain name to access resources?

If you use a CDN acceleration domain name to access resources, factors such as CDN cache may affect the stability of hotlink protection in COS. In this case, you are advised to log in to the [CDN console](#) to configure hotlink protection. For detailed directions, please see [Hotlink Protection Configuration](#).

Can I use an allowlist and also allow accessing the file by opening the URL in a browser?

When setting hotlink protection, you can choose to allow empty referer, so that the file can be accessed by opening its URL in a browser even if an allowlist is configured.

What do I do if I have allowed `a.com` to access the `test` bucket via an allowlist, but the web player for `a.com` cannot play videos stored in the `test` bucket?

When you play videos on websites using players such as Windows Media Player or Flash Player, if the request referer is empty, the allowlist will not be hit. Therefore, you can allow empty referer when configuring the allowlist.

How can I allow accessing COS files only via a corporate network?

You can enable hotlink protection for your bucket and configure a blocklist/allowlist to limit visit sources. Currently, domain names, IPs, wildcard addresses, and more are supported. For more information, please see [Setting Hotlink Protection](#).

Note:

If a CDN acceleration domain name is used for accessing, CDN hotlink protection rules will be executed before COS ones.

If a signature is carried in the access URL or headers, hotlink protection-based verification will not be performed.

What do I do if “You are denied by bucket referer rule” is reported when I access objects via a browser?

This error message indicates that your access is denied by the hotlink protection rules set for your bucket. You can check whether your access complies with the hotlink protection rule. If you access with a browser, empty referer should be allowed. Otherwise, accessing via a browser directly will not be available.

How can I allow only specific IPs to access COS resources?

You can use the hotlink protection feature to configure an IP allowlist. In this case, IPs not included in the allowlist cannot access your COS resources. For detailed directions, please see [Setting Hotlink Protection](#).

Cross-Origin Resource Sharing

Last updated : 2024-12-03 19:46:52

What is CORS and how do I enable it?

Cross-origin resource sharing (CORS) is to request resources over HTTP from a domain for another domain. Two origins that differ in protocol, domain name, or port are treated as different origins. To enable cross-origin access, see [Setting Cross-Origin Access](#) or the best practice documentation [Setting Cross-Origin Access](#).

What should I do if COS denies my access request with headers in the allowlist after I configure CORS?

If your access request is denied, troubleshoot the issue as follows:

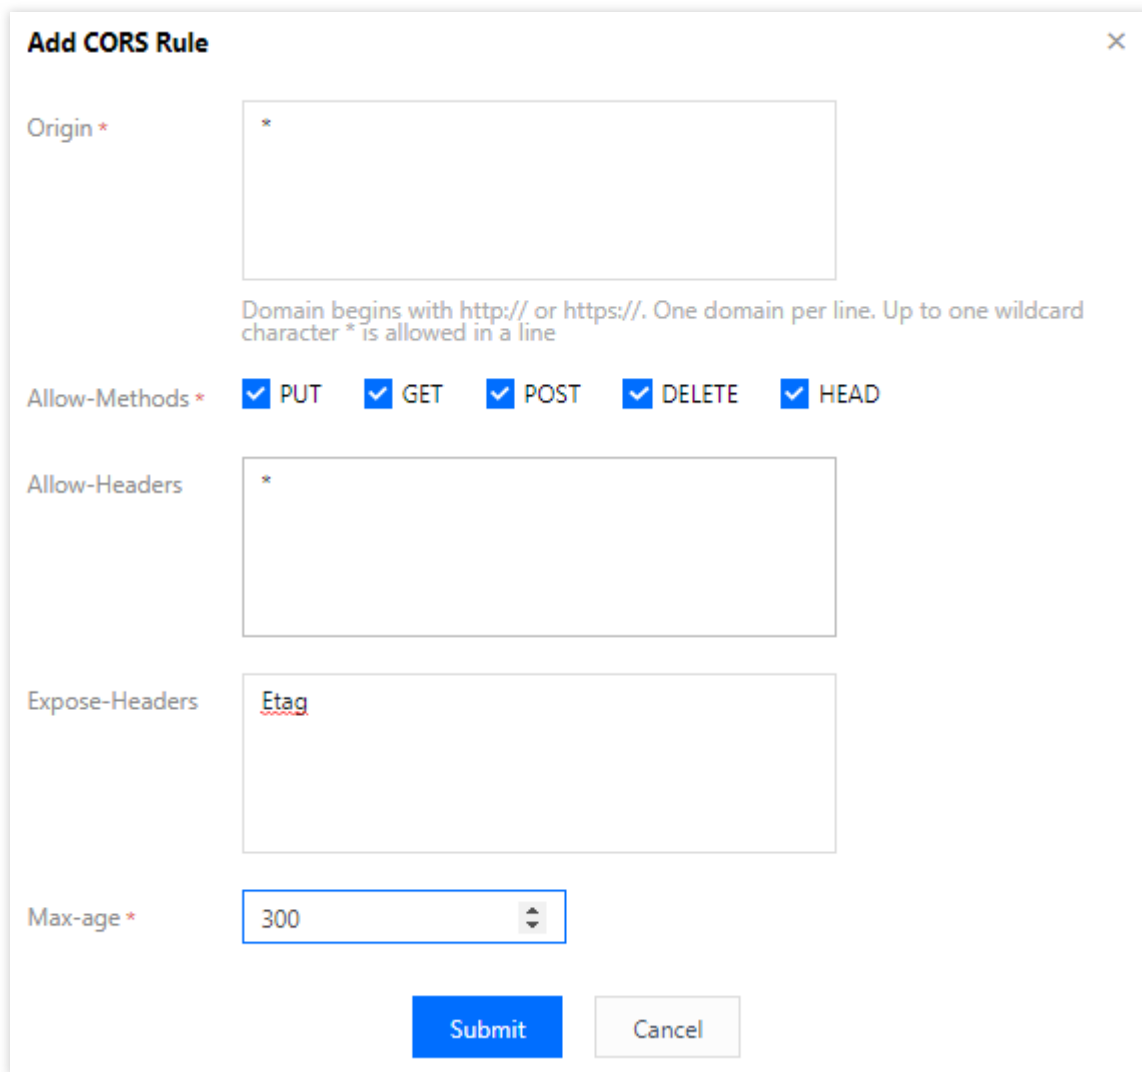
1. Check that your configuration is consistent with the headers your request includes, especially for the presence of invisible characters, such as spaces.
2. Check the domain name with which you initiate the request. If you use a CDN acceleration domain name, you need to configure CORS in the CDN console. For operation details, see [HTTP Response Header](#).
3. Check the permissions of your bucket and determine whether your access will be granted.
4. Check if any error is reported due to your browser's cache. If such an error exists, press **Ctrl + F5** to forcibly refresh your browser, or select **Disable cache** on the **Network** tab of your browser.

How do I configure the file headers in the bucket to return "Access-Control-Allow-Origin:*"?

Set the origin to `*` when configuring CORS. For more information, see the best practice documentation [Setting Cross-Origin Access](#).

What should I do if the error "get ETag error, add `ETag` to CORS ExposeHeader setting." occurs during an upload operation?

Configure the CORS rule as shown below and try using a different browser to test whether it works. For more information, see [Setting Cross-Origin Access](#).



Add CORS Rule

Origin *

Domain begins with http:// or https://. One domain per line. Up to one wildcard character * is allowed in a line

Allow-Methods * ☒ PUT ☒ GET ☒ POST ☒ DELETE ☒ HEAD

Allow-Headers

Expose-Headers

Max-age *

Submit **Cancel**

What should I do if both COS and CDN are used but CORS does not work in COS?

If you are using a CDN acceleration domain name, configure CORS in the CDN console. For operation details, see [HTTP Response Header](#).

Does CORS configuration support fuzzy match of origins?

The console supports fuzzy match of second-level domain names.

What should I do if COS CORS reports an error?

Troubleshoot the issue as follows:

1. Check that CORS rules are configured in the COS console. For operation details, see [Setting Cross-Origin Access](#).
2. Check whether a CDN acceleration domain name is used. If yes, you need to configure CORS rules in the CDN console. For operation details, see [HTTP Response Header](#).
3. If CORS rules are configured, check whether they are effective by running a command in the format of `curl -Ivo /dev/null "<Object address>" -H "origin:<Domain name>"`, for example, `curl -Ivo`

```
/dev/null "https://bucketname-1250000000.cos.ap-guangzhou.myqcloud.com/test.png" -H  
"origin:https://www.baidu.com" . If the status code 200 is returned, the CORS rules are effective. In that  
case, clear the browser cache and try again.
```

4. If the problem persists, try to configure `max-age=0` in the CORS rules.

Can I add IP addresses to CORS rules?

CORS rules support IP addresses. For more information, see [Setting Cross-Origin Resource Sharing \(CORS\)](#).

What should I do if a CORS error is reported when I use CDN to access files in COS although CDN has been configured in COS?

Enable CORS in the CDN console. For operation details, see [HTTP Response Header](#).

What should I do if a CORS error is reported when I access file URLs?

Check whether CORS is configured properly. If yes, we recommend that you clear the browser cache and try again. If the problem persists, try to configure `max-age=0` in the CORS rules. For CORS configuration details, see [Setting Cross-Origin Access](#).

Inventory

Last updated : 2024-12-26 15:17:06

How can I know whether an inventory file has been delivered successfully?

When an inventory file is generated depends on your configuration. If you have selected “daily”, the inventory file will usually be delivered early in the morning on the next day. If you have selected “weekly”, the inventory file will be generated and delivered on the last day of the current week in most cases.

If you want to be notified upon a successful inventory file delivery, you can go to the [SCF console](#) to configure a COS trigger and set **Event Type** to **Use Put Bucket inventory API to create an inventory task**.

How can I analyze an inventory report?

After an inventory report is generated, you can use the [COS Select](#) feature to filter information in the inventory. The following are some examples:

1. Filtering files whose storage class is STANDARD:

```
select * from cosobject s where s._7 = TO_STRING('Standard')
```

2. Filtering files smaller than 5 GB:

```
select * from cosobject s where s._4<5*1024*1024
```

3. Filtering files larger than 5 GB and use the STANDARD storage class:

```
select * from cosobject s where s._4>5*1024*1024 AND s._7=TO_STRING('Standard')
```

4. Filtering files whose status is “replica” (indicating the replication has been completed):

```
select * from cosobject s where s._9=TO_STRING('replica')
```

5. Viewing the first 100 records in the inventory report:

```
select * from cosobject s limit 100
```

How can I export all file information?

You can [enable inventory](#) for your bucket. In this way, COS will regularly (daily/weekly) publish an inventory report that contains the object attributes and configuration details.

Note:

Currently, the inventory feature is not available for Finance Cloud regions.

The inventory feature incurs **Management feature fees**. For detailed pricing, please see [Product Pricing](#).

How can I get a file list?

You can get a file list as follows:

1. Use the COS console to [enable inventory](#) for your bucket. The inventory feature allows you to regularly (daily/weekly) publish inventory reports about the object attributes, configurations, and more. For more information about inventory, please see [Inventory Overview](#).
2. Call the [GET Bucket \(List Objects\)](#) API to obtain a list of all objects. The returned list is in XML format, which can be processed as needed.

Can I reset the inventory configuration immediately after I found out I set it incorrectly?

COS's inventory feature regularly reads the latest configuration at midnight every day before it executes tasks. After you modify the inventory configuration, tasks will be executed early in the morning on the next day.

Does COS support counting file quantity by file type?

You can use the [inventory feature](#) to regularly (daily/weekly) scan specified objects or objects with a specified prefix in a bucket, output an inventory report, and save the CSV file to a specified bucket. After this, you can use "fileFormat" to filter objects by file type and count the quantity.

How can I compare a local file and the one stored in COS?

You can initiate `HEAD Object` or `List Object` requests to obtain the MD5 checksum of one or more objects and compare the value(s) with the local object(s). For large buckets, you can use the [inventory feature](#) to asynchronously obtain the object list as well as the MD5 checksums. For detailed directions, please see [Enabling Inventory](#).

How can I export an XLS file that contains the COS filename, file size, and object URL?

You can enable the inventory feature to automatically output an inventory report, and save the CSV file to a specified bucket. With the inventory feature, you can obtain the file path, file size, last modified time, ETag, storage class, and other information. An object URL can be obtained by combining the bucket's region and the file path. For more information, please see [Inventory Overview](#).

How can I see the number of files in a folder and the storage they take up?

If there are not a lot of files in the folder, you can use the console to view the folder details, which include the number of files as well as the storage they take up. If the number of objects in your bucket is greater than 10,000, you are advised to query using the [inventory feature](#).

Deleting Buckets

Last updated : 2025-03-06 12:06:11

Does It Support Deleting Buckets In Batches?

Unsupported. If needed, refer to [Documentation](#) to delete individually via the console, tool, API, or SDK. Or call the [DELETE Bucket](#) API in a loop.

Is It Possible To Delete Files From a Bucket On a Specified Date?

Not supported. If needed, you can develop to meet your requirement through the API interface. First, [list objects](#), then filter the resources to be deleted by time, and then call the [DELETE Object API](#) to delete.

Version Control Is Enabled, Making It Impossible To Empty the Bucket or Perform Complete Deletions. How Can This Be Resolved?

First, you need to suspend the version control feature of the bucket (for details, see [Setting Versioning](#)), and then [Setting Lifecycle](#) to delete all objects in the bucket (including historical version objects).

Will There Be a Fee If No Files Are Uploaded To an Empty Bucket?

If you do not delete a bucket, even if there are no files in it, normal request fees will still be incurred as long as someone knows the bucket link and sends a request. To save costs, it is recommended to regularly manage and optimize the use of buckets.

Will Deleting a Bucket Incur a Fee?

Yes, the deletion operation is a write request.

How To Handle the Situation Where the Bucket Cannot Be Deleted Due To Arrears?

After arrears exceed 24 hours, COS cannot be read normally. You can top-up at the [Billing Center Console](#) before entering the COS console to delete resources.

Unable To Cancel the Account, Account Resource Check Prompts That the Bucket Needs To Be Deleted, How To Handle It?

This note indicates that you have previously enabled the Cloud Object Storage (COS) service. Please refer to the following steps to proceed:

1. Confirm if there are any buckets in your bucket list that have not been deleted. If the bucket list is not empty, then [all buckets need to be deleted](#).
2. After completing the above steps, please wait for 1 hour before proceeding with the [account cancellation](#) operation. If you still cannot cancel your account, please [submit a ticket](#) for help.

Domain Names and CDN

Static Website

Last updated : 2024-12-26 16:32:33

What should I do if I failed to set a custom domain name in the COS console?

1. Make sure that the domain name has obtained an ICP filing.
2. Make sure that the DNS resolution of the domain name is correct. If CDN acceleration is disabled, you need to go to the DNS console to map the CNAME record of your domain name to the [default domain name](#) of the bucket.

What is the difference between enabling and disabling CDN acceleration if I have my own domain name bound?

CDN acceleration enabled: The domain name is managed by CDN. Enabling CDN acceleration in the COS console has the same effect as adding a domain name in the [CDN console](#) (setting COS as the origin). The CDN-allocated CNAME record is needed for DNS resolution. During the configuration, add the domain name first and then resolve it.

CDN acceleration disabled: The domain name is managed by COS. The domain name configuration is delivered to all download devices connected to the region where the bucket resides. The default domain name of the bucket is used as the CNAME record for DNS resolution.

Why does the `Content-Disposition` header I set for objects not take effect?

Other custom headers can take effect once set. However, `Content-Disposition` takes effect only if the static website hosting feature is enabled and you access objects with a custom domain name.

What should I do if a static website cannot be accessed by using a CDN domain name?

Check the configuration of the CDN-accelerated domain name in the following steps:

1. Select **Static Website Endpoint** as the origin type.
2. Set origin-pull authentication and CDN service authorization based on the bucket permission:
If the bucket permission is private read, authorize the CDN service and enable origin-pull authentication.
If the bucket permission is public read, there is no need to authorize the CDN service or enable origin-pull authentication.

3. Set CDN authentication based on the bucket permission:

(1) If the bucket permission is private read:

| CDN Authentication | Access at CDN Acceleration Domain Name | Access at COS Domain Name | Use Case |
|--------------------|--|-----------------------------|---|
| Disabled (default) | No | COS authentication required | Direct access to the CDN domain name to |

| | | | |
|---------|-----------------------------|-----------------------------|--|
| | | | protect the content on the origin |
| Enabled | URL authentication required | COS authentication required | Full-linkage protection (with hotlink protection for CDN authentication supported) |

(2) If the bucket permission is public read:

| CDN Authentication | Access at CDN Acceleration Domain Name | Access at COS Domain Name | Use Case |
|--------------------|--|---------------------------|--|
| Disabled (default) | Yes | Yes | Site-wide public access via CDN or origin |
| Enabled | URL authentication required | Yes | Hotlink protection enabled for access via CDN but not origin (not recommended) |

4. After confirming that the above configurations are correct, check the protocol used to access the CDN acceleration domain name and the **forced HTTPS** configuration of the static website:

If you are using the HTTP protocol to access the CDN acceleration domain name, **do not enable forced HTTPS**.

If you are using the HTTPS protocol to access the CDN acceleration domain name, we recommend you **enable follow 301/302** for the CDN acceleration domain name. For more information, see [Follow 301/302 Configuration](#).

5. If the problem persists, [contact us](#) for assistance.

What should I do if 404 is returned on refresh when the static website is used together with the frontend Vue framework and the router is set to History mode?

On the static website configuration page of the bucket, set the error document path to the landing page of the web application (generally index.html) and set the status code of the error document to 200. For the configuration directions of static websites, see [Setting Up a Static Website](#).

Note:

After the above configurations are completed, if you need 404 to respond normally, you can configure it at the most bottom layer of the Vue frontend router configuration. In most cases, wildcards match the custom 404 component.

Private Network Domain Name

Last updated : 2024-06-19 15:50:10

Does COS have a private network domain?

COS default origin domain name is in the format of `<BucketName-APPID>.cos.<Region>.myqcloud.com`.

By default, public network and intra-region private network access are supported, for example, `examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com`. For more information on domain name, see [Regions and Access Endpoints](#).

You can use this domain to access COS over a private network (COS will be resolved to a private IP).

In this case, the following will be generated:

Private network traffic: Both private network upstream traffic and downstream traffic are free of charge. For more information, see [Traffic Fees](#).

Requests: The total number of read requests and write requests are calculated every day according to the number of times a request command is sent, with 10,000 times as the minimum unit. For more information, see [Request Fees](#).

Global Acceleration Endpoint

Last updated : 2025-05-13 18:38:11

What are the precautions for using a global acceleration domain name in COS?

Below are precautions for using a global acceleration domain name:

The global acceleration domain name will take effect 15 minutes after being enabled.

After the global acceleration domain name is enabled, the maximum bandwidth for a bucket will be allocated based on the business volume of the entire network.

After the global acceleration domain name is enabled, only requests using that domain name will be accelerated.

However, the default bucket domain name can still be used.

When using a global acceleration domain name, fees will be incurred only for requests for which linkage is accelerated. For example, if you use a global acceleration domain name to upload data from Beijing to a bucket in Beijing, the request will not incur acceleration fees as the linkage was not accelerated.

When using a global acceleration domain name, you can specify to use HTTP or HTTPS transfer protocol. However, if the request is transmitted via a private network Direct Connect line, COS will choose to use HTTPS protocol to guarantee data transfer security.

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

What do I do if the system indicates that the bucket does not exist when I use a global acceleration domain name to access COS via a path containing

`/files/v2/appid/bucketname/` ?

Global acceleration is a feature of COS V5 but `/files/v2/` is a reserved field of COS V4. Using the reserved field in COS V5 will cause an internal logic conflict. Be sure to use a V5 API to access the global acceleration domain name. For more information, please see [API Overview](#).

What operations are currently supported by global acceleration domain names?

Currently, global acceleration domain names support file upload and download. Related APIs are APIs that support acceleration domain names as listed below.

| No. | API | No. | API |
|-----|---------------|-----|-------------------------|
| 1 | PutObject | 7 | ListParts |
| 2 | PostObject | 8 | UploadPart |
| 3 | GetObject | 9 | AbortMultipartUpload |
| 4 | HeadObject | 10 | CompleteMultipartUpload |
| 5 | OptionsObject | 11 | ListMultipartUploads |

| | | | |
|---|-------------------------|---|---|
| 6 | InitiateMultipartUpload | - | - |
|---|-------------------------|---|---|

Under what circumstances will global acceleration be charged?

After global acceleration is enabled, global acceleration traffic will be charged only when a Direct Connect linkage between two Tencent Cloud data centers is used for data transmission, which accelerates the data transmission. For example, if data is uploaded from Xizang to the bucket in Beijing, the Xizang data center will connect to the Chengdu data center first and transmit the data to the storage layer in Beijing through the Direct Connect linkage. In such a case, acceleration fees will be charged. If data is uploaded from Xizang to the bucket in Chengdu, the Xizang data center will directly connect to the Chengdu data center. In this case, the data directly falls on the Chengdu layer, there is no acceleration effect, and there will be no extra charge.

Custom Origin Domain Name

Last updated : 2024-12-26 14:34:54

How do I access objects with my own domain name?

You can do so by binding a custom domain name. For more information, see [Enabling Custom Origin Domain Name](#).

Do I need to obtain an ICP filing from Tencent Cloud if I use a custom domain name?

It depends on the following requirements:

For content delivery in the Chinese mainland, ICP filing is required. You are not required to do so through Tencent Cloud though.

If your domain name is connected to a CDN node outside the Chinese mainland, you don't need to obtain an ICP filing for it.

Does a COS custom domain name support HTTPS?

The feature of configuring HTTPS for custom COS domain names is being upgraded. Currently, certificate hosting is supported in public cloud regions in the Chinese mainland, Singapore, and Silicon Valley, with more regions to come. For unsupported regions, you can configure a reverse proxy for the domain name by referring to [Supporting HTTPS for Custom Endpoints](#).

How does COS return the access links of custom domain names after files are uploaded?

COS currently does not support this feature. However, you can implement it by concatenating access links and using custom domain names to replace default domain names. For more information, see [Regions and Access Endpoints](#).

Do I need to enable CDN if I use custom domain names to access COS?

No. You can log in to the [COS console](#) to set custom domain names. For detailed directions, see [Enabling Custom Origin Domain Name](#).

Why does the original custom domain name disappear from the COS console when the origin is changed in the CDN console?

If you use the COS v5 console and a JSON domain name is configured, the COS v5 console cannot display the new domain name. Check whether the domain name configured in your bucket is in JSON format, and if so, change it to an XML domain name.

Do I need to remove the DNS configuration from the Lighthouse instance before binding a custom domain name to a COS bucket?

Only one CNAME record can be configured for a domain name. Therefore, you need to delete the DNS relationship between the domain name and the Lighthouse instance first and then bind the domain name to the COS bucket.

What should I do if the system prompts that DNS resolution or CNAME is not in effect?

After being configured, DNS resolution or CNAME may take several minutes to take effect. You can wait a while and try to access your bucket at your custom origin domain name again. If the problem persists, log in to your DNS console to check whether the DNS relationship is configured correctly.

CDN Acceleration Domain Names

Last updated : 2025-05-19 09:45:05

How do I enable CDN for COS?

For more information, see [Enabling Custom CDN Acceleration Domain Name](#).

Does COS support CDN origin-pull with HTTPS?

Yes. For more information, see [Setting Origin-Pull](#).

What is the difference between COS and CDN?

COS and CDN are two different products.

Tencent Cloud [Cloud Object Storage \(COS\)](#) is a distributed storage service designed to store a massive number of files. It allows you to upload, download, and manage files in various formats for storage and management of massive amounts of data.

[Content Delivery Network \(CDN\)](#) consists of high-performance cache nodes distributed around the globe to accelerate internet content delivery. These nodes store your content based on caching policies. When a user makes a content request, it will be routed to the node closest to the user, reducing access latency and improving availability.

You do not have to enable CDN when using COS. Using CDN in COS applies to the following scenarios:

Scenarios that require low latency and high download speed

Scenarios where gigabytes to terabytes of data need to be transmitted across regions, countries, and continents

Scenarios where the same content needs to be downloaded frequently and repeatedly

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

Can frontend businesses access the content of COS by means of CDN and temporary key?

No. If you want to use CDN to access COS with private read/write permission on COS, see [CDN origin-pull authentication](#).

Can I access private-read buckets via CDN acceleration?

Yes, but you need to be authorized with related configurations first. For more information, see [Private-read buckets](#).

When a file is updated (re-uploaded or deleted) on COS, its cached content remains unchanged in CDN, resulting in inconsistency with the origin server. Can the cache in CDN be purged automatically when the file on COS is updated?

COS itself does not support automatic purging of CDN cache, a feature you should use with the help of Serverless Cloud Function (SCF). For more information, see [CDN Cache Purging](#).

Can I use COS to upload files via a CDN acceleration domain name?

It is recommended that you do not use a CDN acceleration domain name as a custom domain name to upload files, because CDN itself is not used for accelerated upload. You are advised to use the global acceleration feature of COS to accelerate data upload and download. For more information on global acceleration, see [here](#).

Does COS come with the CDN feature?

No. You need to configure the CDN feature yourself. For more information, see [Enabling Custom CDN Acceleration Domain Name](#).

Other Domain Name Issues

Last updated : 2024-12-26 16:20:58

Does COS support HTTPS access?

Yes. COS allows SSL transport for all access nodes in the [available regions](#), and HTTPS is enabled by default in the SDK and console. **COS strongly recommends you use HTTPS to protect the data linkage for transport. If you use unencrypted HTTP, the linkage may be monitored or the data may be stolen.**

When I manage domain names on the console, I am always prompted to "Enable at least one available key". What should I do?

Log in to the [CAM Console](#) to check if the cloud API key is enabled.

If the cloud API key is not enabled, create and enable a key before managing domain names.

If the cloud API key is enabled and you are still prompted, check whether the account you are operating with is a sub-account (collaborator's or sub-user's account):

If it is a sub-account, log in using the root account and confirm that the cloud API key is enabled.

If it is a root account, purge the browser cache and log in to the Tencent Cloud account again.

What's the difference between COS default domain name, custom CDN acceleration domain name, and custom origin server domain name?

For more information about domain names, please see [Domain Name Management Overview](#).

Default domain name: COS origin server's domain name, which is automatically generated based on the bucket name and region when you create a bucket.

Custom CDN acceleration domain name: you can bind for your bucket a custom domain name to Tencent Cloud Content Delivery Network (CDN), and access objects in your bucket using this domain name. (If you have enabled "Custom domain name" in the legacy COS Console, the new console will continue to display "Custom domain name" instead of "Custom CDN acceleration domain".)

Custom origin server domain name: You can bind a custom domain name with ICP filing to the current bucket, and access objects in the bucket using this domain name.

What's the difference between CDN acceleration and global acceleration in COS?

1. Different use cases: **CDN acceleration** is mainly used for downloading and distributing a great number of objects in a bucket, especially for scenarios where the same data needs to be downloaded repeatedly. For more information, please see [CDN Acceleration Overview](#). **Global acceleration** leverages Tencent's load balancing system that features global traffic scheduling, to intelligently route and parse user requests and select the optimal network linkage for nearby access. COS's global acceleration also speeds up uploads and downloads. For more information, please see [Global Acceleration Overview](#).

2. Different costs. CDN acceleration may involve CDN traffic fees and CDN origin-pull fees, while global acceleration may involve global acceleration traffic fees. For detailed traffic pricing, please see [Traffic Fees](#).

Do I need to bind a domain name to use COS?

No. You can access COS with COS's default domain in the format of `<BucketName>APPID>.cos.<Region>.myqcloud.com` without binding a custom domain name. For more information about COS domain names, please see [Regions and Access Endpoints](#). If you need to bind a custom domain, please see [Enabling Custom Endpoints](#).

Why can I access COS via a private IP range 169.254.0.x in a VPC in the same region?

If the client and COS are in the same region, the COS domain name is forcibly resolved to the IP range 169.254.0.x in the private DNS service of Tencent Cloud through hijacking. By default, the IP range 169.254.0.x communicates with your VPC IP range, and the traffic is diverted to the gateway through internal routes to access COS. Therefore, if you access COS via a private network, do not modify the DNS service configuration. Otherwise, private DNS will fail.

Object Operations

Upload and Download

Last updated : 2024-12-03 17:14:03

Does COS limit the upload and download bandwidth?

15 Gbps of upstream and downstream bandwidth for each bucket in a public cloud region in the Chinese mainland, or 10 Gbps for each bucket in other regions. If this threshold is reached, traffic throttling will be triggered. For more information, see [Specifications and Limits](#).

How can I directly preview a file in my browser without downloading it?

You need to specify a correct `Content-Type` header for this file. In addition, the value of `Content-Disposition` cannot be `attachment`. If the browser supports the current file format, it will directly open the file instead of downloading it. For detailed directions, see [Custom Headers](#).

How do I directly download a file in my browser without previewing it?

You can go to the [COS console](#) and set the value of `Content-Disposition` in the custom object headers to `attachment`. For detailed directions, see [Custom Headers](#).

You can also let your browser pop up a window for the file to be downloaded by setting the value of the request parameter `response-content-disposition` in the GET Object API to `attachment`. For more information, see [GET Object](#).

Note:

To use the `response-*` parameter in a request, the request must be signed.

How do I determine if I am accessing COS over a private network?

Tencent Cloud COS adopts intelligent resolution for COS endpoints. In this way, the optimal linkage can be provided for you to access COS with different ISPs.

If you have deployed a CVM within Tencent Cloud for accessing COS over a private network, you must first ensure that the CVM resides in the same region as the COS bucket, then use the `nslookup` command on the CVM to resolve the COS endpoint. If a private IP is returned, access between the CVM and COS is over a private network; otherwise, it is over a public network.

If your CVM resides in a different region from the COS bucket, but it is still in one of COS regions, you can use the COS private network global acceleration domain name to access files and achieve cross-region access between the CVM and COS.

How to determine access over a private network

Tencent Cloud products within the same region can access each other over a private network, incurring no traffic fees. Therefore, we recommend choosing the same region when you purchase different Tencent Cloud products to save on costs.

Note:

The private networks of Public Cloud regions do not interconnect with those of Finance Cloud regions.

The following shows how to determine access over a private network:

For example, when a CVM accesses COS, to determine whether a private network is used for access, use the

`nslookup` command on the CVM to resolve the COS endpoint. If a private IP is returned, access between the CVM and COS is over a private network; otherwise, it is over a public network.

Note:

Generally, a private IP address takes the form of `10.*.*.*` or `100.*.*.*`, and a VPC IP address takes the form of `169.254.*.*`.

Assume that `examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com` is the address of the destination bucket; the `Address: 10.148.214.13` below indicates access is over a private network.

```
nslookup examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com

Server:                10.138.224.65
Address:                10.138.224.65  #53

Name:   examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com
Address: 10.148.214.13
Name:   examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com
Address: 10.148.214.14
```

For more information on private and public network access and connectivity testing, see [COS Access via Private Network and Public Network](#).

For the private DNS server addresses of CVM, see [Private Network DNS](#).

Note:

The private IPs of Tencent Cloud BM instances may be different from those of CVM instances, and their formats are usually `9.*.*.*` or `10.*.*.*`. If you have any queries, please [contact us](#).

How do I download a folder?

You can log in to the [COSBrowser](#), select the folder to be downloaded, and click **Download** to download the folder or files in batches. You can also download the folder using the COSCMD tool. For more information, see [COSCMD](#).

What should I do if the error "403 Forbidden" occurs or access permission is rejected when I perform upload/download and other operations?

You can troubleshoot by referring to [403 Error for COS Access](#).

How do I upload or download multiple files using COS?

COS allows you to upload or download multiple files through various methods such as the console, APIs/SDKs, and tools.

Console: For detailed directions, see [Uploading Objects](#) and [Downloading Objects](#).

APIs/SDKs: COS allows you to operate on multiple files programmatically by repeatedly calling an API or SDK. For more information, see [APIs for object uploads/downloads](#) and [SDK Overview](#).

Tools: You can use tools such as [COSBrowser](#), [COSCMD](#), and [COSCLI](#) for batch operations.

When I upload a new file to a bucket in which another file of the same name exists, will the old file be overwritten or will the new file be saved with a different version name?

The versioning feature is now available in COS. If versioning is not enabled for the bucket, when you upload a new file to a bucket in which another file of the same name already exists, the older one will be directly overwritten; if versioning is enabled, multiple versions of the object will co-exist.

What is the minimum part size of a multipart upload in COS?

1 MB. For more information, see [Specifications and Limits](#).

When uploading large files using multipart upload, can I replace an invalid signature to continue the multipart upload?

Yes.

How do I generate a temporary URL for files in COS?

For more information, see [Download via Pre-Signed URL](#).

I have set a validity period for a signature, but why can it still be used to download objects after it has expired?

By default, the browser will cache objects that have been loaded successfully. Therefore, if you access the same URL, the cached object will be returned without requesting the server again. Therefore, you are advised to use the

`Cache-Control: no-cache` header during object upload to prevent browser caching (see [PUT Object](#) or [Initiate Multipart Upload](#) for details). Alternatively, you can specify the `response-cache-control=no-cache` request header during object download to prevent browser caching (see [GET Object](#) for details).

What should I do if I upload a file on the console and "Failed to upload. System error." is displayed?

This error may occur due to an unstable local network environment. You can try the upload again in a different network environment.

How do I prevent others from downloading my COS files?

To do so, you can set your bucket permission to private read/write. For more information, see [Setting Access Permission](#). You can also configure a hotlink protection whitelist on your bucket to block any access from domain names not in the list. For more information, see [Setting Hotlink Protection](#).

Can I use case-insensitive download URLs?

No. COS filenames are case-sensitive, and thus so are the download URLs. If you have enabled CDN acceleration for your bucket, you can go to the CDN console to configure [Cache Ignore URL Case](#), which will increase the hit rate to some extent.

What should I do if the error "your policy or acl has reached the limit (Status Code: 400; Error Code: PolicyFull)" occurs when I am uploading files or creating a bucket?

COS allows each root account to have up to 1,000 bucket ACLs. If more bucket ACLs have been configured, this error will occur. Therefore, you can delete unnecessary bucket ACLs.

Note:

You are not advised to use object-level ACLs or policies. When calling APIs or SDKs, if you do not need ACL control over a file, we recommend leaving the ACL-related parameters (such as x-cos-acl and ACL) empty to inherit the bucket permissions.

Data Management

Last updated : 2024-12-03 17:27:46

FAQs About Deletion

What should I do if buckets fail to be deleted via the console with a message "The directory is not empty" or "Please delete the valid data in the bucket"?

1. Verify whether your console version is V4 or V5. If you're using version V4, [Contact Us](#) to apply for upgrading to V5.
2. Log in to the [COS Console V5](#), go to the details page of the bucket to be deleted, and click **Incomplete Upload** to delete the file fragments.
3. Return to the **Bucket List** to delete the bucket.

Can I restore a file that has been deleted by mistake in my bucket?

Currently, files that users delete manually by mistake cannot be restored. However, you can enable versioning on your bucket to help recover future data lost due to accidental deletion or application failure. It is because versioning allows you to upload and store multiple versions of an object with the same name in your bucket, so that you can extract, delete, or restore a specified version of the object. For more information, see [Setting Versioning](#).

FAQs about Incomplete Multipart Upload

How is an incomplete multipart upload generated, and can I download parts uploaded incompletely?

An incomplete multipart upload is generated when you suspend or cancel an object upload. COS allows you to upload a large file greater than 5 GB in multiple file parts. During a multipart upload job, uploaded parts will be stored as an incomplete multipart upload and cannot be downloaded if you don't call the Abort Multipart Upload or Complete Multipart Upload API.

Will the uploaded parts of an incomplete multipart upload take up my storage capacity and incur fees?

Like objects, incomplete multipart uploads consume your storage capacity, and incur storage usage fees.

How do I (regularly) clear incomplete multipart uploads?

You can delete an incomplete multipart upload directly by using the COS console. For directions, see [Deleting Incomplete Multipart Upload](#). Alternatively, you can [regularly clear incomplete multipart uploads by configuring a lifecycle](#).

Will clearing incomplete multipart uploads affect other complete multipart uploads?

No, it won't. It just deletes the unsuccessful uploads.

Custom Headers

Can object headers be customized in batches?

COS supports custom headers. For more information, see [Custom Headers](#).

FAQs About Other Features

Does COS support setting callbacks? For example, creating a thumbnail for each image uploaded and save it to another bucket?

You can realize this by using COS with SCF. For more information, see [Acquire Image on COS and Create a Thumbnail](#).

Does COS allow me to view the size of a folder?

COS allows you to view the number and size of objects in the current folder. For more information, see [Viewing Folder Details](#).

Can I set a COS object back to its last version?

Once you enable [Versioning](#) on your bucket, you can store multiple versions of an object in the bucket, and extract, delete or recover a specified version of the object. For detailed directions, see [Setting Versioning](#).

How do I view the number of files of a certain type (e.g. image) in COS?

To do this, you can enable inventory, and check your generated inventory report. For more information, please see [Enabling Inventory](#).

Storage Class

Last updated : 2024-12-26 16:35:08

How is INTELLIGENT TIERING billed?

INTELLIGENT TIERING fees include **INTELLIGENT TIERING storage usage fees** and **INTELLIGENT TIERING object monitoring fees**.

1. INTELLIGENT TIERING storage usage fees are charged differently depending on the storage tier of objects.
- When objects are stored in the frequent access tier, STANDARD storage usage fees are charged.
- When objects are stored in the infrequent access tier, STANDARD_IA storage usage fees are charged.

Note:

STANDARD and STANDARD_IA storage usage fees vary by region. For details about their pricing, see [Product Pricing](#).

Request fees and traffic fees are also incurred during object upload and download. For the billing examples, see [Traffic Fees](#) and [Request Fees](#).

2. INTELLIGENT TIERING object monitoring fees are charged based on the number of objects stored (excluding files smaller than 64 KB). For details about their pricing, see [Product Pricing](#).

Sample

Assume that an organization has 100,000 objects (10 MB per object, 1 TB in total), and the data is stored in the INTELLIGENT TIERING storage class in Beijing region. If 20% of the objects (i.e., 20,000 objects) transition to the infrequent access tier every month, the object monitoring fees and storage usage fees for each month will be as follows:

| Month | Object Monitoring Fees (USD) | INTELLIGENT TIERING Storage Usage Fees (USD) | STANDARD Storage Usage Fees (USD) |
|-------|------------------------------|---|-----------------------------------|
| 1 | 0.25 | $1024 \times 0.024 = 24.58$ | $1024 \times 0.024 = 24.58$ |
| 2 | 0.25 | $819.2 \times 0.024 + 204.8 \times 0.018 = 23.35$ | $1024 \times 0.024 = 24.58$ |
| 3 | 0.25 | $655.36 \times 0.024 + 368.64 \times 0.018 = 22.36$ | $1024 \times 0.024 = 24.58$ |
| 4 | 0.25 | $524.288 \times 0.024 + 499.712 \times 0.018 = 21.58$ | $1024 \times 0.024 = 24.58$ |
| 5 | 0.25 | $419.4304 \times 0.024 + 604.5696 \times 0.018 = 20.95$ | $1024 \times 0.024 = 24.58$ |
| 6 | 0.25 | $335.54432 \times 0.024 + 688.45568 \times 0.018 = 20.45$ | $1024 \times 0.024 = 24.58$ |

As you can see, using INTELLIGENT TIERING reduces storage costs over time. Only a small amount of monitoring fees is incurred each month.

What types of objects is INTELLIGENT TIERING suitable for?

INTELLIGENT TIERING is suitable for large objects (such as audios, videos, and logs) whose access patterns change. Larger average object sizes mean that you pay less for monitoring per GB of objects. If your business has relatively fixed data access patterns, you can set lifecycle configuration to specify the time to transition objects to STANDARD_IA without the need to use INTELLIGENT TIERING.

How do I store objects in INTELLIGENT TIERING?

You can store objects in INTELLIGENT TIERING as follows:

Incremental objects: You just need to specify the storage class as INTELLIGENT TIERING when uploading objects.

Existing objects: You can modify the storage class to INTELLIGENT TIERING by using the `COPY` API. You can also use the lifecycle feature to transition STANDARD or STANDARD_IA objects to INTELLIGENT TIERING.

Note:

INTELLIGENT TIERING objects smaller than 64 KB will always be stored in STANDARD. For such objects, we recommend you upload them to the STANDARD, STANDARD_IA, ARCHIVE, or DEEP ARCHIVE storage class as needed to reduce costs.

How do I disable INTELLIGENT TIERING configuration?

INTELLIGENT TIERING configuration **can't be disabled** once enabled. If you don't need to store your objects in INTELLIGENT TIERING, you just need to specify the storage class as a non-INTELLIGENT TIERING class such as STANDARD, STANDARD_IA, ARCHIVE, or DEEP ARCHIVE when uploading objects.

Logging and Monitoring

Last updated : 2024-12-26 16:33:30

Logging

Does COS provide file upload, download, and deletion logs?

COS provides the [logging](#) feature, which records the access details of a source bucket. These logs are then stored in a destination bucket for better bucket management. To get the file upload, download, and deletion logs, enable access logging to record file operations.

How do I query which files incur the most public network traffic in COS?

You can use the [logging](#) feature of COS to download bucket access logs and write a program to analyze which files consume the most public network traffic. You can also load the logs to Data Lake Compute (DLC) for statistics collection.

How do I query from which source IPs most public network traffic in COS comes from?

You can use the [logging](#) feature of COS to download bucket access logs and write a program to analyze from which source IPs most public network traffic comes from. You can also load the logs to DLC for statistics collection.

Can I set thresholds for public network downstream traffic and request count in COS?

You can [set an alarm policy](#) in the [CM console](#) to receive alarm notifications when the public network downstream traffic in COS reaches the threshold. COS currently can't automatically suspend the service when the threshold is reached.

How do I view file deletion logs?

You can query the logs shipped by the [logging](#) feature to view file deletion logs. After access logging is enabled, you can load log files to DLC to filter deletion logs. Below is a sample deletion log. You can search for the `DELETE` operation in the `reqMethod` field to get such logs:

```
1.0 examplebucket-1250000000 ap-chengdu 2020-02-10T13:07:00Z examplebucket-1250000000
```

If you cannot find deletion logs among access logs, check whether rules of deletion upon expiration are set in the [lifecycle configuration](#).

How do I query COS bucket configuration logs?

Bucket configuration logs are shipped to CloudAudit. You can search for such logs as instructed in [Viewing Event Details in Operation Record](#).

Where can I query bucket creation and deletion logs?

Bucket creation and deletion logs are shipped to CloudAudit. You can select `DeleteBucket` and `PutBucket` events to filter operation logs as instructed in [Viewing Event Details in Operation Record](#).

Monitoring

Can traffic be throttled in COS?

No. However, you can [create an alarm policy](#) in **CM** to trigger alarms and push notifications by email or SMS when the traffic reaches a certain threshold.

Why does the request count or traffic on the monitoring dashboard increase suddenly?

If your business has an abnormal surge in the request count or traffic, your business may be hotlinked. You need to check whether public read is enabled for your bucket. We recommend you not enable public read, as it will bring uncontrollable risks to your business. You can grant access according to the [principle of least privilege](#).

If you must use public read, we recommend you use the following methods to guarantee the bucket security:

1. Enable the [logging](#) feature for your bucket to log bucket access requests.
2. Enable the [hotlink protection](#) feature to block access requests from abnormal IPs.
3. [Create a COS alarm policy](#) and set a threshold, so that alarm notifications will be sent to you by SMS or email once the traffic exceeds the threshold.

Permission Management

Last updated : 2024-12-26 16:23:36

Keys

How can I view the key information such as `APPID` , `SecretId` , and `SecretKey` ?

The second half of a bucket name is the `APPID` . You can view it by logging in to the [COS console](#). To view information such as `SecretId` and `SecretKey` , log in to the CAM console and go to the [Manage API Key](#) page.

How long will a temporary key be valid?

Currently, a temporary key can be valid for up to 2 hours (i.e., 7,200 seconds) for the root account, and 36 hours (i.e., 129,600 seconds) for a sub-account. The default validity period is 30 minutes (i.e., 1,800 seconds). Requests carrying an expired temporary key will be denied. For more information, please see [Generating and Using Temporary Keys](#).

What do I do if my key information such as `SecretId` and `SecretKey` is compromised?

You can delete the compromised key and create a new one. For more information, see [Access Key](#).

How can I generate a time-bound access URL for a Private Read/Write file?

You can set the effective time for your key by referring to [Generating and Using Temporary Keys](#).

Permissions

How can I grant a sub-account permission to access a specific folder?

You can grant such permission by referring to [Setting Folder Permissions](#). To grant more advanced permissions to a sub-account, see [Authorization Cases](#).

What do I do if COS returns a 403 error?

You can use the [diagnostic tool](#) provided by the COS team, which uses `RequestId` to troubleshoot errors.

1. Check whether the configuration of `BucketName` , `APPID` , `Region` , `SecretId` , and `SecretKey` is correct.
2. If the above information is correct, check whether a sub-account is used. If so, check whether the sub-account has been authorized by the root account. If it has not been authorized, log in to the root account to authorize the sub-account.
3. Perform authorization. For more information, see [Cases of Permission Setting](#).

4. If a temporary key is used for operation, check whether the current operation is in the Policy set when obtaining the temporary key. Otherwise, modify the relevant Policy settings.

What do I do if “AccessDenied” is reported?

In most cases, this error is reported due to unauthorized access or insufficient permissions. You can troubleshoot as follows:

1. Check whether the configuration of `BucketName` , `APPID` , `Region` , `SecretId` , and `SecretKey` is correct. Note that you should also check whether spaces are carried.
2. If the above information is correct, check whether a sub-account is used for operation, and if yes, check whether the sub-account has been authorized by the root account. If it hasn't yet, log in using the root account to authorize the sub-account. For more information on authorization, see [Cases of Permission Settings](#).
3. If a temporary key is used, check whether the current operation is in the policy set when the temporary key is obtained; if not, modify the relevant policy settings. For more information, please see [Generating and Using Temporary Keys](#).

The COS team provides the [diagnostic tool](#), which uses `RequestId` to troubleshoot errors.

What do I do if the number of bucket permissions has reached the upper limit?

Each root account (i.e., each `APPID`) can have up to 1,000 bucket ACLs. If more bucket ACLs have been configured, an error will be reported. Therefore, you are advised to delete unnecessary ACLs.

Note:

You are not advised to use file-level ACLs or policies. When calling APIs or SDKs, if you do not need ACL control over a file, we recommend leaving the ACL-related parameters (such as `x-cos-acl` and `ACL`) empty to inherit the bucket permissions.

What should I do if an error is reported during bucket creation?

Possible causes:

1. A bucket with the same name already exists. In this case, you need to name the bucket differently.
2. The public read/private write or public read/write permissions have been set for too many existing buckets, and the maximum number of ACL rules for the root account has been reached. When you create a bucket, an error will be reported because this maximum number cannot be adjusted.

The following two solutions are provided for your reference:

Solution 1: Set the access permission of an existing bucket to private read/write and try creating a bucket again. For more information, see [Setting Access Permission](#).

Solution 2: In **Permission Policy Settings**, click **Add Policy** and configure access permissions as needed. For more information, please see [Adding Bucket Policies](#).

Can I access a public-read file using a signed URL whose signature has expired?

If you use an expired signed URL to access a public-read file, COS will first verify the permissions. If the URL has expired, the access will be denied.

What do I do if "403 Forbidden" or "permission rejected" is reported during uploads, downloads, or other operations?

You can troubleshoot as follows:

1. Check whether the configuration of `BucketName` , `APPID` , `Region` , `SecretId` , and `SecretKey` is correct.
2. If the configuration above is correct, check whether a sub-account is used for the operation. If yes, check whether the sub-account has been authorized by the root account. If it has not yet been authorized, log in using the root account to authorize the sub-account. For more information about authorization, please see [Cases of Permission Settings](#).
3. If a temporary key is used, check whether the current operation is in the policy set when the temporary key is obtained; if not, modify the relevant policy settings. For more information, please see [Generating and Using Temporary Keys](#).

The COS team provides the [diagnostic tool](#), which uses `RequestId` to troubleshoot errors.

How can I prevent users from downloading COS data?

You can prevent users from downloading data based on your use case as follows:

1. To prevent sub-accounts from downloading data, please see [Granting Sub-accounts Access to COS](#).
2. To prevent anonymous users from downloading data, you can set your bucket to private-read/write, or set `deny` `anyone Get Object` in the bucket policy.

How can I grant permissions to a sub-account under another root account?

Assume that you (root account A) need to grant bucket permissions to the sub-account B0 that is under the root account B. You need to first grant the root account B permissions to operate your bucket. Then, root account B should grant the permissions to its sub-account B0. For detailed directions, see [Granting Bucket Permissions to a Sub-Account that is Under Another Root Account](#).

How can I only allow sub-accounts/collaborators to upload but not delete files?

You can log in to the [CAM console](#) to create a custom policy that grants specified permissions to sub-accounts. For detailed directions, please see [Creating Custom Policies](#).

Note:

When creating the custom policy, you need to grant read permissions, set upload only for write operations, and **do not grant delete permissions**.

Visual Policy Generator JSON

▼ COS(0 actions)

Effect * ☒ Allow ☐ Deny

Service * COS (cos)

Action *

Select actions

☐ All actions (cos:*)

Action Type

☒ Read (35 selected) ▶

☐ Write ▼

Select Action

Filter Actions 🔍

| <input type="checkbox"/> | Action Name | Description |
|--------------------------|--------------------|---------------------|
| <input type="checkbox"/> | AbortMultipartU... | |
| <input type="checkbox"/> | AppendObject | |
| <input type="checkbox"/> | CompleteMultip... | |
| <input type="checkbox"/> | Create/Job | Create a COS Bat... |
| <input type="checkbox"/> | DeleteBucket | |
| <input type="checkbox"/> | DeleteBucketCORS | DeleteBucketCORS |

Support for holding shift key down for multiple selection

When I access a public-read bucket using its default domain name, how can I hide the returned file list?

You can set a permission to deny anyone's Get Bucket operation for the bucket by following the steps below:
Log in to the [COS console](#), click **Bucket List**, click the desired bucket, and select **Permission Management**.

Method 1:

1. Click **Permission Policy Settings**. Then, click **Add Policy** under **Visual Editor**.
2. Configure the permissions as shown below and then click **OK**.

Add Policy

When dealing with authorizations, it is recommended that you strictly comply to principles of least privilege. You can authorize the user to perform restricted operations (such as only authorize read operations) and access only the resources with specified prefix, to avoid data security risks due to excessive permissions and operations that you don't mean to authorize.

Effect *
☐ Allow
☒ Deny

User *

| User Type | Account ID | Operation |
|-----------|------------|-----------|
| Add User | | |

Resource *
☒ The whole bucket
☐ Specific resources

Resource path *
examplebucket-1250000000/*

Operation *

| Action Name | Operation |
|-------------|-----------|
| GetBucket | Delete |
| Add Action | |

Condition

| Name | Operator | Value ⓘ | Operation |
|---------------|----------|---------|-----------|
| Add Condition | | | |

OK
Cancel

Method 2:

Click **Permission Policy Settings**. Then, click **JSON > Edit** and enter the following code:

```
{
  "Statement": [
    {
      "Action": [
        "name/cos:GetBucket",
        "name/cos:GetBucketObjectVersions"
      ],
      "Effect": "Deny",
      "Principal": {
        "qcs": [
          "qcs::cam::anyone:anyone"
        ]
      },
      "Resource": [
        "qcs::cos:ap-beijing:uid/1250000000:examplebucket-1250000000/*"
      ]
    }
  ]
}
```

```
    }  
  ],  
  "version": "2.0"  
}
```

Note:

Replace values in `qcs::cos:ap-beijing:uid/1250000000:examplebucket-1250000000/*` with the actual information as follows:

“ap-beijing”: Replace it with the region where your bucket resides.

“1250000000”: Replace it with your `APPID` .

“examplebucket-1250000000”: Replace it with your bucket name.

The second half of the bucket name is the `APPID` . You can view it by logging in to the [COS console](#).

Are COS's ACLs bucket-specific or account-specific? Can I specify permissions when uploading files?

ACLs are account-specific. You are not advised to use file-level ACLs or policies. When calling APIs or SDKs, if you do not need ACL control over a file, we recommend leaving the ACL-related parameters (such as `x-cos-acl` and `ACL`) empty to inherit the bucket permissions.

How do I authorize a collaborator to access a specified bucket?

A collaborator is a special sub-account. For more information, please see [Access Policy Language Overview](#).

Can I isolate permissions by buckets or other dimensions if I have multiple services that need to work with buckets?

You can log in to the [CAM console](#) to go to the **User Management** page, where you can create sub-accounts for different services and grant different permissions for them.

How can I create sub-accounts for subsidiaries or employees and grant them permissions to access specific buckets?

You can create sub-accounts and grant them permissions by referring to [Granting Sub-accounts Access to COS](#).

How can I allow specific sub-accounts to only operate a certain bucket?

To grant a sub-account access to a specific bucket, you can add a bucket policy for the sub-account with the root account. For more information, see [Adding Bucket Policies](#).

Ranger Authentication and Verification

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

Other

What do I do if I cannot access COS resources normally?

You can troubleshoot by referring to [Resource Access Error](#).

What do I do if “HTTP ERROR 403” is returned when I access COS using a CDN domain?

This is usually because the CDN acceleration domain is disabled. You can troubleshoot by referring to [“HTTP ERROR 403” Is Returned When I Access COS Using a CDN Domain](#).

What do I do if I use a CDN domain name to access COS but only access the old version of the file?

This is usually because of the existing cache. You can troubleshoot by referring to [A URL Points to a Wrong File](#).

Can the front-end access COS using CDN and a temporary key?

If you need authentication when CDN pulls from COS with files set to private read/write, see [Setting CDN Acceleration](#).

Data Processing Function Service

Last updated : 2024-12-03 19:46:26

Does COS support file decompression?

File decompression is a data processing solution provided by Tencent Cloud COS based on [Serverless Cloud Function \(SCF\)](#). For more information, see [File Decompression](#).

Does COS file decompression decompress compressed files in second-level directories?

No. However, you can adjust the function logic to implement the feature.

Does COS support automatic compression upon file upload?

No.

Does COS support automatic CDN purging?

You can configure automatic CDN purging via [Serverless Cloud Function \(SCF\)](#). For operation details, see [CDN Cache Purging](#).

Can I back up cloud database data to COS?

You can configure the database backup feature via [Serverless Cloud Function \(SCF\)](#). After you configure backup function rules for the specified bucket, SCF will regularly scan for your database backup files and dump them to the bucket. For operation details, see [Setting Cloud Database Backup](#).

Data Security

Last updated : 2024-12-03 17:15:00

Versioning

Can I restore accidentally deleted data?

No. However, you can enable versioning for your bucket so that you can store multiple versions of an object in a bucket, and extract, delete, or restore a specific object version. Versioning allows you to restore data lost due to accidental deletion or application failures. For more information, see [Setting Versioning](#).

What is COS's solution to disaster recovery?

You can achieve disaster recovery in COS by:

1. Enabling [versioning](#), which allows you to store multiple versions of an object in the bucket. For detailed directions, see [Versioning Configuration](#).
2. Using [cross-bucket replication](#) to achieve remote disaster recovery. For more information, see [Cross-Bucket Replication Configuration](#).
3. Using the [MAZ configuration](#), which provides IDC-level disaster recovery capabilities for your data.

Note:

1. Currently, the MAZ configuration of COS is supported only in Guangzhou, Shanghai, and Beijing regions and will be available in other public cloud regions in the future.
2. Using the MAZ configuration incurs high storage usage fees. For more information, see [Pricing | Cloud Object Storage](#).

How can I delete noncurrent object versions after I enable versioning for a bucket?

You can [set a lifecycle rule](#) and enable **Managing historical versions** to transition or delete noncurrent object versions.

Can a newly uploaded object not overwrite the old one that has the same name?

No. By default, the old object with the same name will be overwritten by the new one. However, you can enable [versioning](#) for your bucket so that multiple object versions can be preserved. For more information, see [Versioning Overview](#).

How can I download a specific version of object?

If you download with APIs or SDKs, add the `versionId` request parameter. For the API calling directions, see [GET Object](#).

If you download via the console, set the historical versions to **Display** in the top navigation bar so that you can download the desired object version.

How can I delete noncurrent object versions in batches?

You can use the COSBrowser tool to one-click delete noncurrent object versions in batches. For more information, see [COSBrowser User Guide for Desktop Version](#).

You can also [configure a lifecycle policy](#) to delete objects that were modified more than 1 day ago for noncurrent object versions.

Cross-Region Replication

Does cross-region replication use a private or public network?

By default, cross-region replication uses a private network.

Note:

Note that cross-region replication incurs traffic fees, which cannot be redeemed with a resource pack yet. The fees incurred will be deducted from your account at 00:00 the next day.

Can I sync resources between two regions?

Yes. Resources under the same account can be synced between two regions. You can [set cross-bucket replication](#) to replicate objects incrementally.

How to quickly migrate resources from one account to another account?

You can use [COS Migration](#) to migrate data between buckets. Alternatively, you can [set cross-bucket replication](#).

Does cross-bucket replication support replicating existing data?

No. You can use [Batch Operation](#) instead.

If I enable cross-bucket replication and delete a file from the source bucket, will the file be deleted as well in the destination bucket?

In a source bucket with cross-bucket replication enabled, COS will replicate the following:

Any new objects uploaded to the source bucket after the cross-bucket replication rule is added.

Object attributes such as object metadata and version ID.

Object operations, such as adding an object of the same name (equivalent to adding a new object) and deleting an object.

Note:

If you specify an object version to delete in the source bucket by specifying a version ID, COS will not replicate this delete operation.

If you add a bucket-level configuration such as a lifecycle rule to the source bucket, COS will not replicate any resulting object operations.

For more information, see [Cross-Bucket Replication Actions](#).

Data Encryption

Does COS support data encryption?

Yes. COS supports encryption such as bucket encryption and object encryption. For more information, see [Setting Bucket Encryption](#) and [Setting Object Encryption](#).

Does COS encryption affect performance?

A client-side/COS-managed/KMS key is used to encrypt the file content into ciphertext, which affects performance to some extent (mainly by increasing access delay). The delay does not significantly affect large object reads/writes, but has a certain impact on small object reads/writes.

How can I get an encrypted object?

To get an encrypted object, include an encryption header when reading it. The encryption header differs according to the encryption algorithm. For more information, see [Common Request Headers](#).

Content Security

Why are non-compliant files found in my COS bucket?

Your data is stored in COS, and the data access permission is public read. When you access and disseminate such data on the public network, you need to comply with applicable laws and regulations. If the content of such data violates regulations, the compliance team of Tencent Cloud will handle it accordingly, and handled files will be displayed in the **Violation List** in the COS console.

I have already enabled the content moderation feature, but why do I still receive violation notifications?

Possible reasons:

1. The content moderation feature is not configured correctly; for example:

Automatic freezing is not configured, or the identified non-compliant data is not processed in time (such as deleting files).

The configured data freezing score is too high, so some non-compliant files have low scores and are not frozen.

Some non-compliant images are historical data and have not been moderated. We recommend you conduct a full moderation of historical data to check the entire bucket.

2. If the moderation configuration is correct, but non-compliant data is determined to be normal, this is generally because the data is relatively obscure, and the existing moderation model doesn't correctly moderate it. We will regularly collect similar moderation error samples for continuous optimization. You can also [submit a ticket](#) for customized moderation services.

Other

Are there backups for the STANDARD, STANDARD_IA, and ARCHIVE storage classes?

COS data is stored at the underlying layer using multiple replicas or erasure coding (both are imperceptible to users). The storage engines are distributed across multiple availability zones in a region, making the data reliability 99.999999999%.

SDKs

Last updated : 2024-12-26 16:36:13

Common SDK Issues

What do I do if I upload a file using a file stream/handle, but the uploaded file is truncated or the size is 0?

If the file is uploaded with a file stream/handle, the stream or handle usually contains an offset pointer. If the stream/handle is used before being uploaded, the offset pointer may not point to the starting point of the file. In this case, the SDK uploads the file starting from where the offset points to by default, resulting in the truncated or 0-size file. To solve this problem, you are advised to check the offset, or point the offset to the file's starting point manually if necessary.

How do I get the object's URL after it is successfully uploaded?

URLs of objects in COS are formed using a fixed format. For more information, please see [Object Overview - Object access address](#).

What do I do if the system reports that the temporary key expired when I upload files?

Please follow the steps below for troubleshooting:

1. Check whether the time of the machine that runs your applications is correct. If the machine time is incorrect, rectify it.
2. Check whether the expiration time (`expirationDate`) that you set is earlier than the current time. The current time being later than the expiration time will cause the signature to expire. In that case, you need to change the expiration time and regenerate a temporary key.
3. The iOS SDK uses the `QCloudSignatureProvider` and `QCloudCredentailFenceQueueDelegate` protocols during initialization. The `QCloudCredentailFenceQueue` scaffolding tool will cache and reuse your temporary key. You can update the cache by reinitializing the `credentialFenceQueue` instance to avoid using a temporary key that has expired. For more information, see [Create a COS instance](#).

How do I determine whether files are successfully uploaded?

In COS, each object corresponds to an `ETag` value. After a file is successfully uploaded, an `ETag` value of the String type will be returned and the `ETag` value is not NULL. You can add a determination condition to determine whether files are successfully uploaded.

How do I request an object for which hotlink protection is configured?

Add a Header with a specified referer to your request for an object.

Does generating a pre-signed URL generate network requests and incur fees? Will there be a delay?

Generating a pre-signed URL is local logic and does not generate network requests, causing no additional network latency and no additional cost. You can call the SDK API to generate a pre-signed URL at any time when needed.

How do I configure a custom domain name for a COS pre-signed URL?

A pre-signed URL contains a fixed default domain name, which can be replaced via encoding.

How do I create a directory in COS SDK?

A directory in COS is virtual and is actually an object ending with `/`. You can call the object upload API to create an object ending with `/`, which is a directory. For more information, see [Mini Program SDK use case: Create a directory](#).

Why are different results returned when I use the same prefix rule and data structure to obtain `ObjectList` via COS SDK?

To make it easier for you to get started, COS simulates the display mode of "folder" or "directory" in the **console and graphical tools such as COS browser**. This is realized by creating an empty object with a key value of `project/` and displaying it as a traditional folder. Therefore, the `objectList` obtained through the SDK will contain empty objects whose object names end with `/`.

SDK FAQs

[Android SDK FAQs](#)

[C SDK FAQs](#)

[.NET \(C#\) SDK FAQs](#)

[Go SDK FAQs](#)

[iOS SDK FAQs](#)

[Java SDK FAQs](#)

[JavaScript SDK FAQs](#)

[PHP SDK FAQs](#)

[Python SDK FAQs](#)

[Mini Program SDK FAQs](#)

Tools

COSBrowser

Last updated : 2024-12-03 17:52:45

What is COSBrowser?

COSBrowser is a visual interface tool launched by Tencent Cloud COS to make it easier and simpler for you to view, transfer, manage, and interact with COS resources. Currently, it is available for desktop (Windows, macOS, Linux) and mobile clients (Android, iOS). For more information, see [COSBrowser Overview](#).

How do I download COSBrowser?

For the download URL and instructions, see [COSBrowser Overview](#).

How do I log in to COSBrowser?

For detailed directions, see [User Guide for Desktop Version](#) or [Mobile Version Features](#).

Desktop Version login

COSBrowser Desktop Version can be logged in only with a TencentCloud API key.

Parameter description:

1. TencentCloud API **secretID** and **secretKey** : You can get them on the [API Key](#) page in the CAM console. After you log in successfully, the key will be saved in **Historical Key** for future use.
2. Bucket/Access Path: You can leave it empty when logging in with the root account. If you use a sub-account for login, you need to enter an authorized path such as `example-1250000000/test/` .

Note:

You cannot log in to COSBrowser with a project key.

Mobile Version login

COSBrowser Mobile Version supports the following three login options:

Login with WeChat: If your Tencent Cloud account was created through WeChat or associated with a specific WeChat account, you can use the WeChat account to quickly log in to COSBrowser.

Login with email: If your Tencent Cloud account was created through email or associated with a specific email address, you can log in to COSBrowser by entering the email address and password.

Login with permanent key: You can log in using your TencentCloud API key (SecretId and SecretKey; project key is not supported), which can be obtained on the [API Key](#) page in the CAM console. After successful login, the account will be kept logged in permanently.

Note:

If your Tencent Cloud account was created with a QQ account, you can also use the login with WeChat method to log in just by selecting login with QQ on the redirected WeChat Mini Program screen.

If you use a sub-account, you can log in with key or WeChat. For login with WeChat, just select the sub-account on the redirected WeChat Mini Program screen.

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

Why can't I find the storage path when I log in to COSBrowser with a sub-account?

1. Make sure that the sub-account has permission to access COS. For more information, see [Granting Sub-accounts Access to COS](#).
2. If the sub-account only has permission to access a specified bucket or a specified directory under a bucket, you need log in to COSBrowser and manually add a storage path in "Bucket" or "Bucket/Object-prefix" format (for example, `examplebucket-1250000000`), and select the region where the bucket should reside.

Key Login [Advanced Setting >](#)

SecretID

AKID7wXsRnJIAI8G7

SecretKey

.....

Buckets/Access Path ⓘ

examplebucket-1250000000/

Region

ap-chengdu ▼

Remark

Not required, Add remark

Login

[Get SecretKey](#) [History key](#) [Local Logs](#)

Can I log in to COSBrowser with a temporary key?

Login with a temporary key is not supported.

How do I enter the trial version of COSBrowser?

Notes for trial

Application trial rules:

After entering the trial version of the application, COSBrowser will automatically generate a temporary account and log in. The temporary account is for one-time use. After exit, it will be automatically logged out of, with all data erased.

The temporary account is valid for 24 hours. If you want to continue the trial after expiration, click again on this page.

Application trial restrictions:

The trial version only provides basic data management capabilities, such as uploading files, downloading files, and sharing links. To try out more features, log in with your personal account. For more information, see [Getting Started with COSBrowser](#).

What should I do if I can't launch the COSBrowser client by double-clicking on the CentOS GUI?

You can run the `./cosbrowser.AppImage --no-sandbox` command in the terminal to launch COSBrowser.

What are the system requirements for COSBrowser installation?

Currently, COSBrowser are available on Desktop Version and Mobile Version.

Desktop Version

Requirements for Windows: Windows 7 32/64-bit or later or Windows Server 2008 R2 64-bit or later

Requirements for macOS: macOS 10.13 or later

Requirements for Linux: GUI- and AppImage-enabled distributions

Mobile Version

Requirements for Android: Android 4.4 or later

Requirements for iOS: iOS 11 or later

For the download address, see [Download URL](#).

What is the file sync feature of COSBrowser?

You can use the **file sync** feature of COSBrowser Desktop Version to upload specified files in your local folders to a bucket in real time. For detailed directions, see the description of the file sync feature in [User Guide for Desktop Version](#).

Can I see all file thumbnails at a time in the file list in COSBrowser?

COSBrowser currently can't directly display the thumbnails of all files.

Why are only three buckets displayed in the list on COSBrowser Mobile Version?

The overview page of COSBrowser Mobile Version displays three buckets by default. You can scroll down to view more buckets.

Can I use COSBrowser to directly upload objects to the STANDARD_IA storage class?

COSBrowser uploads objects to the STANDARD storage class by default. You can select the storage class and access permission when uploading objects.

How do I transfer a large number of files at a greater speed?

Take COSBrowser for Windows as an example. You can navigate to **Advanced Settings**, and configure a greater number of concurrent files and parts under **Upload** and **Download**.

How do I copy a file link in COSBrowser?

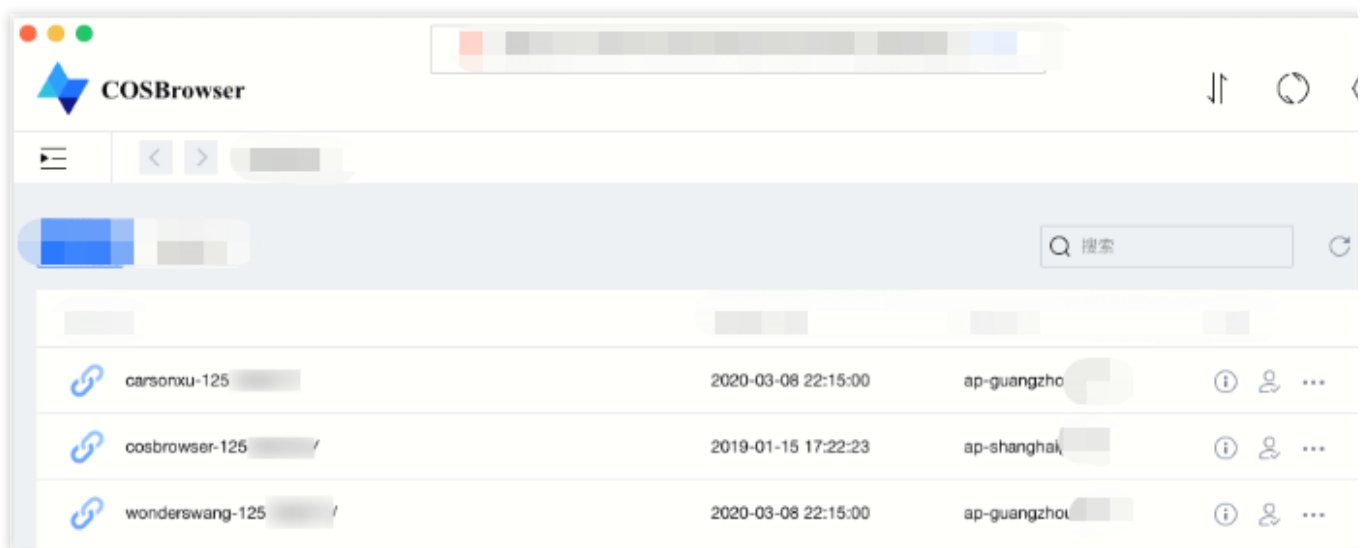
You can copy a file link as follows:

1. Select the target file in the file list and right-click **Copy Link** to open the **Copy Custom Link** window.
2. In the file list, click **Details** to open the **Details** window, directly copy the **object address** or create a **temporary link**.

Note:

If public read is enabled for the file, you can use an unsigned link, i.e., object address (valid permanently) to access it. If private read is enabled for the file, you must use a signed link to access it. You can customize the link validity period in the **Copy Link** window, which is 2 hours by default.

What should I do if the error "Update failed due to permission denied" is reported when I run COSBrowser on macOS?



Cause

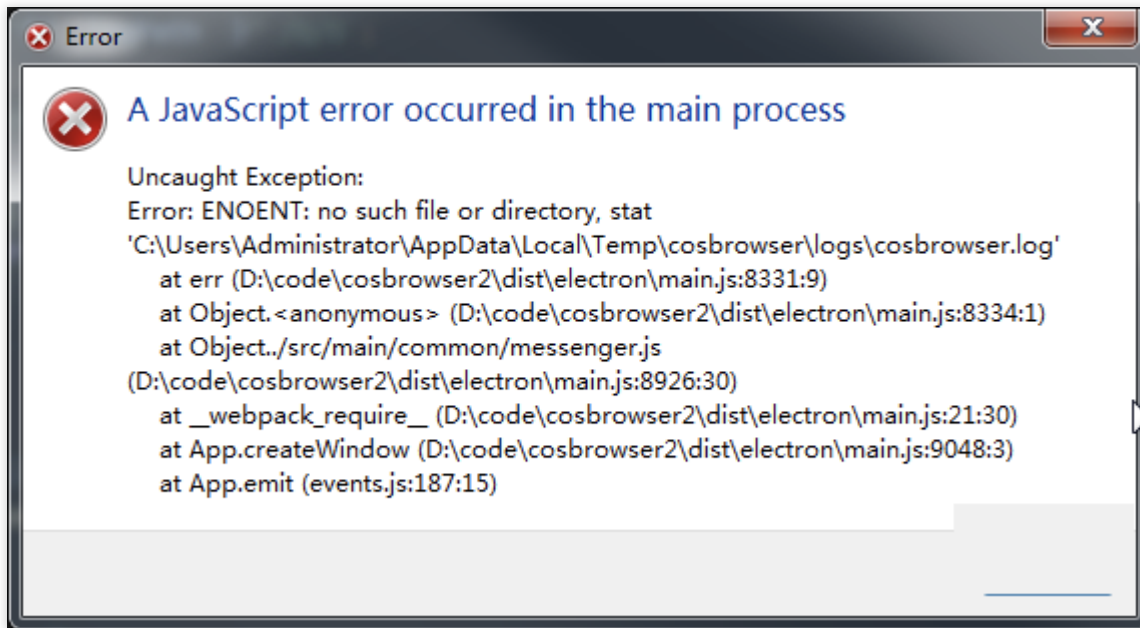
The above issue may occur when the two files `com.tencent.cosbrowser` and `com.tencent.cosbrowser.ShipIt` under the `/Users/username/Library/Caches/` directory belong to a root account and a user account, respectively. The resulting permission issue can cause update failures.

Solution

Run the following command in Terminal on your Mac:

```
sudo chown $USER ~/Library/Caches/com.tencent.cosbrowser.ShipIt/
```

What should I do if the error "no such file or directory, stat 'C:\\Users\\XXX\\AppData\\Local\\Temp\\cosbrowser\\logs\\cosbrowser.log'" is reported and COSBrowser cannot be used?



Solution: We recommend you download v2.1.x or later.

What should I do if `cosbrowser.exe` installation is interrupted?

Cause

This problem occurs if COSBrowser has been installed before, but its traces have not been removed from your system after manual deletion. Therefore, when you reinstall it, the installer aborts the installation as it identifies existing traces but no actual application.

Solution

Manually remove the traces left after the previous installation of COSBrowser, or use a cleaner tool (such as the Software Management tool provided by Tencent PC Manager).

What should I do if a DNS error is reported when I go to the file list in COSBrowser?

The DNS error indicates that the COS domain name failed to be resolved in your local network. We recommend you change your local DNS server address to a public one such as `114.114.114.114` and try again or change the network environment for testing.

COS Migration

Last updated : 2024-12-03 17:18:02

What should I do if the migration tool exits abnormally?

The migration tool supports checkpoint restart. If the upload of a large file is interrupted due to the tool error or service failure, you can run the tool again and restart the upload from the checkpoint.

If the files that have been migrated successfully to COS are deleted through the console or other methods, will the migration tool upload them again?

No. All successfully uploaded files will be recorded in the `db` directory. Before the migration tool runs, the `db` directory will be scanned first, and the files recorded in it will not be uploaded again. For more information, see [Migration mechanism and process](#).

What should I do if the migration fails with the message "403 Access Deny" displayed in the log?

Check whether your key, bucket, and region information is correct and ensure that you have the operation permission. If you're using a sub-account, it needs to be authorized by the root account. If you migrate data from a local file system or other cloud storage, write/read permission on the bucket is required. For the "Bucket copy" operation, read permission on the source bucket is also required.

What should I do if the migration to COS from another cloud storage fails with an error message "Read timed out"?

This error occurs when the data download from other cloud storage times out due to insufficient bandwidth. For example, when you migrate overseas data from AWS to COS, "read time out" may occur due to network latency caused by insufficient bandwidth. To solve this problem, you can increase the network bandwidth and test download speed with `wget` before migration.

What should I do if the migration fails with a message "503 Slow Down" shown in the log?

This error occurs when frequency control is triggered. A limit of 30,000 QPS is imposed on an account in COS. We recommend you decrease the concurrency for small files in configuration. Then run the tool again to resume the migration.

What should I do if the migration fails with the message "404 NoSuchBucket" shown in the log?

Check whether your key, bucket, and region information is correct.

What should I do if an exception occurs with the following message?

```
Exception in thread "main" java.lang.ExceptionInInitializerError
    at org.rocksdb.RocksDB.loadLibrary(RocksDB.java:64)
    at org.rocksdb.RocksDB.<clinit>(RocksDB.java:35)
    at org.rocksdb.Options.<clinit>(Options.java:25)
    at com.qcloud.cos_migrate_tool.record.RecordDb.init(RecordDb.java:43)
    at com.qcloud.cos_migrate_tool.task.TaskExecutor.initRecord(TaskExecutor.java:94)
    at com.qcloud.cos_migrate_tool.task.TaskExecutor.run(TaskExecutor.java:146)
    at com.qcloud.cos_migrate_tool.app.App.main(App.java:60)
Caused by: java.lang.UnsupportedOperationException: Cannot determine JNI library name for ARCH='x86' OS='windows 10'
e='rocksdb'
    at org.rocksdb.util.Environment.getJniLibraryName(Environment.java:78)
    at org.rocksdb.NativeLibraryLoader.<clinit>(NativeLibraryLoader.java:19)
    ... 7 more
```

The possible reason is that the tool uses RocksDB which requires 64-bit JDK. Check whether your JDK version is x64.

What should I do if the message "rocksdb's jni library cannot be found" is displayed in a Windows environment?

In a Windows environment, the tool needs to be compiled in Microsoft Visual Studio 2015. In case of the above error message, you need to install [Visual C++ Redistributable for Visual Studio 2015](#).

How do I modify the log level?

Modify the file `src/main/resources/log4j.properties` by replacing the value of `log4j.rootLogger` with the log level, such as `DEBUG`, `INFO`, and `ERROR`.

What do I do if the `/tmp/librocksdbjnixxx.so: ELF file OS ABI invalid` error is reported in the Linux environment?

IFUNC needs to be supported on Linux and the binutils version in the running environment should be later than 2.20.

What should I do if a task fails to be fully executed and "java.nio.file.FileSystemLoopException" is reported in `error.log` ?

The exception information in `error.log` is similar to:

```
2022-XX-XX XX:XX:XX [ERROR] [main:xxx]
[com.qcloud.cos_migrate_tool.task.MigrateLocalTaskExecutor:]
[MigrateLocalTaskExecutor.java:183]
walk file tree error
java.nio.file.FileSystemLoopException: /dataseal/xx1/file1
at java.nio.file.FileTreeWalker.visit(FileTreeWalker.java:294)
at java.nio.file.FileTreeWalker.next(FileTreeWalker.java:372)
at java.nio.file.Files.walkFileTree(Files.java:2706)
at
com.qcloud.cos_migrate_tool.task.MigrateLocalTaskExecutor.buildTask(MigrateLoca
```

```
lTaskExecutor.java:176)
at com.qcloud.cos_migrate_tool.task.TaskExecutor.run(TaskExecutor.java:244)
at com.qcloud.cos_migrate_tool.app.App.main(App.java:135)
```

The reason is that the file `"/dataseal/xx1/file1"` to be migrated may be a soft link pointing to a resource in its parent directory. You can check with the following command:

```
[root@TENCENT64 /dataseal/cos_migrate_tool_v5-master/log]# ll
/dataseal/xx1/file1
lrwxrwxrwx 1 xx xx xx xx  x xxxx /dataseal/xx1/file1 -> ../xx1/
```

As shown above, the soft link file `"/dataseal/xx1/file1"` points to `"/dataseal/xx1/"` in its parent directory, which causes an infinite loop in traversal. Therefore, the migration task is automatically terminated.

We recommend you delete such files in advance (note: the method of excluding such files in the configuration item `excludes` is invalid).

If other problems occur, try to run the migration tool again. If the problem persists, compress the configuration (with the key hidden) as well as the log directory and [contact us](#).

COSCMD

Last updated : 2024-12-03 17:26:36

What should I do when COSCMD is unavailable?

Please check whether the following requirements are met:

1. The OS is Windows, Linux, or macOS.
2. Local characters use UTF-8 encoding. Otherwise, exceptions will occur when you operate on Chinese files.
3. The local time is in sync with UTC. If there is a large deviation between the two, COSCMD might not function properly.

For more information, please see [COSCMD](#).

Does COSCMD support regular expressions?

No.

I can successfully create a bucket with a name containing uppercase letters using COSCMD, but when I perform other operations with a bucket name containing uppercase letters, an error occurs. What is the reason for this?

COSCMD automatically converts uppercase letters to lowercase ones. A bucket name can only contain lowercase letters, numbers, and hyphens (-), with a length not greater than 50 characters. For more information, see [Specifications and Limits](#).

Can the files in a sub-directory be excluded when I upload or download the files in the root directory with COSCMD?

Yes. You need to use the `--ignore /folder/*` parameter.

For example, if you want to exclude a folder during download, use `coscmd download --ignore /folder/*` to filter out the files in the folder. If you want to ignore files with a certain suffix in the folder, be sure to append `,` to the `"*"` character, or enclose it with `" "`.

How can I transfer a large number of files with quicker speed?

Configure an appropriate value for `MAX_THREAD`, which defaults to 5. The number of threads depends on the server performance, and generally, setting its value to 30 can easily take up full bandwidth. For example, you can set the number of concurrent threads to 30 by running the following command:

```
coscmd config -m 30
```

Does COSCMD support using * to determine objects with a specified prefix to download?

No. You need to use the following command format for download:

```
coscmd download prefix/ localpath/ -r
```

Can I use the `list` command to list files by file upload time in COSCMD?

No. You can list files by specifying a prefix. For more information, please see [COSCMD](#).

Can I use COSCMD to manage buckets of different accounts at the same time?

You can configure only the bucket under one account in the `cos.conf` configuration file. If you need to manage a bucket under another account, switch to the bucket first by using the following command:

```
coscmd config -a SecretID -s SecretKey -b BucketName-APPID -r region
```

`SecretID` and `SecretKey` can be obtained in the [CAM console](#). `BucketName-APPID` indicates the bucket name, and `region` indicates the region where the bucket resides.

Can I specify multiple buckets in the COSCMD configuration file?

You can configure only one bucket in the COSCMD configuration file. If you need to manage another bucket, specify the bucket name and region COSCMD command for bucket switching.

Use the `-b <bucketname-appid>` parameter to specify the bucket name, which must be formatted as `BucketName-APPID`.

Use the `-r <region>` parameter to specify the region where the bucket resides.

Does COSCMD verify filename duplication for uploaded files?

No. If you upload a file whose name is duplicated with that of an existing file, COSCMD will overwrite the existing file.

How do I transfer a large number of files with a quicker speed?

Configure an appropriate value for `MAX_THREAD`, which defaults to 5. The number of threads depends on the server performance, and generally, setting its value to 30 can easily take up full bandwidth. For example, you can set the number of concurrent threads to 30 by running the following command:

```
coscmd config -m 30
```

Does COSCMD verify the content of uploaded files?

No. COSCMD adopts the overwrite upload mode by default. If you need to skip existing identical files, add the `-rs` parameter.

How do I skip existing identical files when uploading files with COSCMD?

You can use the `-rs` parameter to skip files with the same MD5 checksum. For more information, please see [Uploading a folder](#) in [COSCMD](#).

How do I skip identical files when downloading files with COSCMD?

When downloading a folder, you can use the `-s` or `--sync` parameter to skip identical files that already exist locally (provided that the downloaded files were uploaded via the COSCMD `upload` API and the `x-cos-meta-md5` header was included). An example of the complete command is as follows: `coscmd download -rs --skipmd5 cos_path local_path .`

Can I upload multiple folders at the same time with COSCMD?

No. You can upload only one folder at a time. You can put multiple folders that need to be uploaded into a single folder for uploading, but it takes time to copy files locally.

FTP Server

Last updated : 2024-12-03 17:27:29

How do I enable the FTP feature?

COS is a persistent storage that supports web-based requests but does not provide native FTP access. Intermediate transfer is required to use the FTP protocol. **It is recommended that you set up your service by using the [FTP Server Tool](#) provided by Tencent Cloud.**

As an outdated protocol, FTP protocol is unable to verify data integrity, ensure transfer security, or be integrated with the CAM system. Therefore, it is not recommended to use the FTP protocol for access, and Tencent Cloud will provide support for the FTP protocol and intermediate transfer software.

For data synchronization, it is recommended to use the [COS Migration tool](#) or the [COSCMD tool](#).

What does the `masquerade_address` option in the configuration file do and when does it need to be configured?

The `masquerade_address` is a server address configured for the client. When the FTP server runs on a host that is mapped to an external IP through NAT, you need to configure the `masquerade_address` option as an FTP server external IP that the client can access in order to notify the client to use the IP for data communication with the server. For example, assume that you execute `ifconfig` on the host where the FTP server is running, and get a private ENI IP `10.xxx.xxx.xxx`, which is mapped to the public IP `119.xxx.xxx.xxx`. At this time, if the FTP Server does not explicitly set `masquerade_address` to the public IP (119.xxx.xxx.xxx) that the client uses to access the server, the FTP Server in Passive mode may use the private IP (10.xxx.xxx.xxx) to return packets to the client. As a result, the client is able to connect to the FTP Server, but cannot return data packets to the client properly.

After the `masquerade_address` option is correctly configured, I can log in to the FTP server normally, but when I run the FTP command for fetching data such as "list" or "get", the error "The server returns a non-routable address" or "ftp: connect: No route to host" occurs. How do I deal with it?

The most possible reason for this is that the FTP server's `iptables` or firewall policy is configured to reject or drop all ICMP protocol packets. After receiving the data connection IP returned by the FTP server in the PASSIVE mode, the FTP client will send an ICMP packet first to verify the connectivity of the IP. In this case, errors such as "The server returns a non-routable address" may occur.

Solution: Configure the `iptables` policy to only reject or drop the ICMP packet types you want to block. For example, if you only want to block the external ICMP packets of Ping type, you can change the policy to: `iptables -A INPUT -p icmp --icmp-type 8 -s 0/0 -j [REJECT/DROP]`.

Alternatively, you can allow the IP of the client that will access the FTP server.

Why is the uploaded part retained in COS when the upload of a large file is canceled halfway?

The FTP server for the latest version of COS provides a complete streaming upload feature. When you upload a large file, the cancellation or disconnection of the upload will trigger the completion of upload. In this case, COS considers that your data stream has been uploaded and combines the uploaded data into a complete file. If you want to resume the upload, you can upload the file with the same name to overwrite the original one, or delete the incomplete file manually and upload the file again.

What will happen if the size of an uploaded file exceeds the limit?

If the size of the uploaded file exceeds the limit set in the configuration file, the system returns an `IOError` exception and marks the error message in the log.

If you have any other questions, please [contact us](#) and provide the complete `cos_v5.log` log to facilitate troubleshooting.

Why does a limit on the size of file to be uploaded need to be set in the COS FTP server configuration?

For a multipart upload in COS, the maximum number of file parts to be uploaded is 10,000, and the size of each file part is limited to 1 MB to 5 GB. The purpose of imposing the limits is to reasonably calculate the size of a file part. The FTP server supports uploading a single file less than 200 GB by default. But it is not recommended to set the limit to a too large value, because a larger file size limit will cause a larger buffer for file parts during upload, thus increasing the consumption of your memory. Therefore, you are advised to set a reasonable file size limit as needed.

Hadoop

Last updated : 2024-12-03 17:52:16

Use/Consultation

What is Hadoop-COS?

Hadoop-COS is a tool that helps integrate big-data computing frameworks including Apache Hadoop, Spark, and Tez. It allows you to read and write Tencent Cloud COS data just as you do with HDFS. It can also be used as Deep Storage for Druid and other query and analysis engines.

How do I use the Hadoop-COS jar file for self-built Hadoop?

Change the Hadoop-COS POM file to keep its version the same as that of Hadoop before compilation. Next, put the Hadoop-COS jar and COS JAVA SDK jar files in the directory `hadoop/share/hadoop/common/lib`. For more information, see [Hadoop-COS](#).

Is there a recycle bin mechanism in the Hadoop-COS tool?

The recycle bin feature of HDFS is not applicable to COS. When you use Hadoop-COS to delete COS data by running the `hdfs fs` command, the data will be moved to the `cosn://user/${user.name}/.Trash` directory, but no actual deletion will occur, so the data will still remain in COS. You can use the `-skipTrash` parameter to skip the recycle bin feature and delete the data directly. To implement periodic data deletion like with the HDFS recycle bin, configure a lifecycle rule for objects prefixed with `/user/${user.name}/.Trash/`. For the configuration guide, see [Setting Lifecycle](#).

CosFileSystem Class Not Found

Why do I receive the following message during loading, prompting me that the class CosFileSystem was not found: “Error: java.lang.RuntimeException: java.lang.ClassNotFoundException: Class org.apache.hadoop.fs.CosFileSystem not found”?

Possible cause 1

The configuration was loaded correctly, but the `hadoop classpath` does not include the location of Hadoop-COS jar.

Solution

Load the location of Hadoop-COS jar to `hadoop classpath`.

Possible cause 2 `mapreduce.application.classpath` in the `mapred-site.xml` configuration file does not include the location of Hadoop-COS jar.

Solution

Add the path of cosn jar to `mapreduce.application.classpath` in the `mapred-site.xml` configuration file, and restart the service.

The screenshot shows the Ambari web interface. On the left, the 'MapReduce2' service is selected under the 'HDFS' category. The main panel displays the 'Advanced' configuration tab for the 'Advanced mapred-site' group. The 'mapreduce.application.classpath' property is visible, with its value containing the path to the cosn jar file: `.../usr/hdp/current/hadoop-libs/hadoop-hdfs/hadoop-hdfs-2.6.0-cosn.jar`. A red arrow points to the 'cosn' part of the path, indicating where the path should be added.

Why am I receiving a prompt that the class `CosFileSystem` was not found when I use Apache Hadoop?

COS offers two versions: Apache Hadoop and Hadoop-COS, which differ in the configuration of `fs.cosn.impl` and `fs.AbstractFileSystem.cosn.impl`.

Apache Hadoop:

```
<property>
  <name>fs.cosn.impl</name>
  <value>org.apache.hadoop.fs.cosn.CosNFileSystem</value>
```

```
</property>
<property>
  <name>fs.AbstractFileSystem.cosn.impl</name>
  <value>org.apache.hadoop.fs.cosn.CosN</value>
</property>
```

Tencent COS:

```
<property>
  <name>fs.cosn.impl</name>
  <value>org.apache.hadoop.fs.CosFileSystem</value>
</property>
<property>
  <name>fs.AbstractFileSystem.cosn.impl</name>
  <value>org.apache.hadoop.fs.CosN</value>
</property>
```

Frequency Control and Bandwidth

Why am I receiving a 503 error?

In big data scenarios, high concurrency may trigger the COS frequency control policy, resulting in a 503 `Reduce your request rate` error exception. You can configure retries for your failed requests by configuring the `fs.cosn.maxRetries` parameter, which defaults to 200 retries.

Why hasn't my bandwidth limit setting gone into effect?

The bandwidth limit setting `fs.cosn.traffic.limit (b/s)` is supported only by the latest versions of Hadoop-COS with Tag 5.8.3 or above. For more information, please see [Github](#).

Parts

How do I set a reasonable block (part) size for multipart uploads via Hadoop-COS?

Hadoop-COS uploads large files to COS via concurrent uploads of multiple parts. You can control the size of each part by configuring `fs.cosn.upload.part.size (Byte)`.

Because a COS multipart upload allows at most 10,000 parts for a single file, you need to estimate the largest possible file size you may need to upload to determine the value of this parameter. For example, with a part size of 8 MB, you can upload a single file of up to 78 GB in size. A maximum part size of 2 GB is supported, meaning that the largest single file size supported is 19 TB. A 400 error will be thrown if the number of parts exceeds 10,000. If you encounter said error, please check if you have configured this parameter correctly.

Why can't I see a large file immediately after it was uploaded to COS?

Hadoop-COS uploads all large files greater than the block size (`fs.cosn.upload.part.size`) through multipart upload. You can see the file on COS only after all of its parts were uploaded. Currently, Hadoop-COS does not support Append operations.

Buffers

Which buffer type should I choose for my uploads? What's the difference between them?

You can choose a buffer type by setting `fs.cosn.upload.buffer` to one of the following three values:

`mapped_disk`: default. You need to put `fs_cosn.tmp.dir` under a directory large enough to avoid a full disk in runtime.

`direct_memory`: uses JVM off-heap memory (out of JVM control; not recommended)

`non_direct_memory`: uses JVM on-heap memory; set to 128 MB (recommended).

Why do I get the following buffer creation failure when I set the buffer type to

`mapped_disk` : create buffer failed. buffer type: mapped_disk, buffer factory:org.apache.hadoop.fs.buffer.CosNMappedBufferFactory?

Possible cause

You do not have the read or write permission on the temporary directory used by Hadoop-COS. The directory is `/tmp/hadoop_cos` by default, and can be customized by configuring `fs.cosn.tmp.dir`.

Solution

Obtain the read and write permission on the temporary directory used by Hadoop-COS.

Runtime Exceptions

What should I do if the following exception is thrown when I perform computing tasks:

`java.net.ConnectException: Cannot assign requested address (connect failed) (state=42000,code=40000)?`

Generally, when this exception occurs, you have established too many short TCP connections in a short period of time. After the connections are closed, local ports will enter a 60-second timeout period by default instead of being immediately repossessed. As a result, there is no available port during this period for your Client to establish a socket connection with the Server.

Solution

Modify the `/etc/sysctl.conf` file with changes to the following kernel parameters:

```
net.ipv4.tcp_timestamps = 1          #Enables support for TCP timestamp
```

```
net.ipv4.tcp_tw_reuse = 1          #Supports the use of a socket in the status of TIME
net.ipv4.tcp_tw_recycle = 1       #Enables quick repossession of a socket in the stat
net.ipv4.tcp_syncookies=1        #Enables SYN Cookies. The default value is 0. When
net.ipv4.tcp_fin_timeout = 10     #Waiting time after the port is released
net.ipv4.tcp_keepalive_time = 1200 #The interval for TCP to send KeepAliv
net.ipv4.ip_local_port_range = 1024 65000 #The range of ports for external conne
net.ipv4.tcp_max_tw_buckets = 10240 #The maximum number (default: 180000)
```

When I upload a file, why does the exception "java.lang.Thread.State: TIME_WAITING (parking)" occur with "org.apache.hadoop.fs.BufferPoll.getBuffer" and "java.util.concurrent.locks.LockedBlockingQueue.poll" locked in the stack?

Possible cause

You may have initialized the buffer repeatedly, but not actually triggered the write action.

Solution

Change the configuration to the following:

```
<property>
  <name>fs.cosn.upload.buffer</name>
  <value>mapped_disk</value>
</property>
<property>
  <name>fs.cosn.upload.buffer.size</name>
  <value>-1</value>
</property>
```

APIs

Last updated : 2024-12-03 17:12:45

Do COS APIs support the S3 protocol?

COS provides APIs that are compatible with AWS S3. For more information, please see [Accessing COS Using the AWS S3 SDK](#).

What do I do if an error message such as "Request has expired" is displayed when I call an API?

There are two possible causes:

The signature has expired when you initiate the request.

Your local system time is out of sync with the local time in your time zone.

For the former case, you are advised to get a new valid request signature before using the API. For the latter case, you need to sync your system time with the time in your time zone.

How do I call an API to delete an object that is not uploaded completely?

First, call the `ListMultipartUploads` API to list the multipart uploads. Then, call `Abort Multipart upload` to abort the multipart upload and delete the uploaded parts.

What do I do if a success response is returned for the batch deletion API, while the object failed to be deleted?

Check the object path, which should not start with a `/`.

What do I do if "NoSuchUpload" is returned for the `UploadPart` request?

If the values of `uploadId` and `partNumber` are the same, newly uploaded parts will overwrite previous ones. If `uploadId` does not exist, "404 NoSuchUpload" will be returned. For more information, please see [Upload Part](#).

How do I modify the storage class for an object using an API?

You can call the `PUT Object - Copy` API to modify `x-cos-storage-class`. For more information, please see [PUT Object - Copy](#).

How do I set the COS signature to be permanent?

A timestamp is used to determine whether the COS signature has expired and it cannot be set to permanent. If you use a permanent key to generate the signature and want the signature to be permanent, you can set the timestamp to be a long time (for example, 50 years) after the current time. If your signature is generated using a temporary key, which is valid for only up to 2 hours, your signature will also be valid for only 2 hours.

Can I call an API to query the COS bill?

No. You can go to the console and view [Bill Details](#). To call APIs to view the bill details, please see [DescribeBillDetail](#).

Can I call an API to query the size of an object?

Yes. You can call [GET Bucket \(List Objects\)](#) to query the size of an object.

Can I call an API to modify the name of an object?

You can call [PUT Object - Copy](#) to copy the object and specify a name for the replicas.

Can I call an API to query the bucket domain name?

You can call the [HEAD Bucket](#) API to query the bucket domain name. The `x-cos-bucket-region` parameter in the response header indicates the region where the bucket resides.

Can I call an API to query the bucket size?

COS does not provide an API to query the bucket size. You are advised to use [Cloud Monitor APIs](#) to query the storage by storage class, and then sum up the storage of each storage class to get the bucket size.

How can I call APIs to query the usage details?

You can:

1. Use the [API request tool](#).

Does COS provide APIs to operate directories?

Technically speaking, COS does not have directories or folders. In fact, folders displayed in the console are empty objects whose names end with a slash (/).

How can I create a directory/folder using APIs?

You can call the [PUT Object](#) API and upload an empty object whose name ends with a slash (/).

Note:

COS does not have the concepts of directories or folders. To make COS more user-friendly, COS simulates folders/directories in GUIs such as the console and COSBrowser. You can create an empty object whose name ends with a slash (/), and it will be displayed as a folder.

How can I call APIs to delete a directory/folder?

COS APIs support only deleting a single file. If you need to delete the entire directory, you can call the [GET Bucket \(List Objects\)](#) API to obtain the list of all objects that have the same specified prefix. Then, call [DELETE Object](#) to delete them.

How to know on which access tier my INTELLIGENT TIERING object is stored?

You can use `x-cos-storage-tier` returned for the [HEAD Object](#) API to query the access tier of your object.

How can I call APIs to search for an object?

You can call [HEAD Object](#) to determine whether the object exists. To search for a specific object, call [Get Bucket](#) to query all objects in the bucket and find your object.

When I call the `GET Object` API, can I download the dynamically specified content that is returned as an attachment?

When you call `GET Object`, carry the **response-content-disposition** parameter in the URL and set its value to **attachment**. Note that this kind of `GET Object` requests require a signature carried. You can generate a signature using COS's signature generation tool.

What do I do if “NoSuchKey” is returned for the `putObjectCopy` request?

Check whether the file exists. If yes, the error is usually caused because the slash (/) is missing from the folder's key. You can add the slash and then retry.

Can I call an API to query how many times an object is requested?

COS does not provide such an API. However, you can [Set Logging](#) and then obtain the requested times by analyzing the logs.