

Cloud Object Storage

FAQs

Product Documentation



Copyright Notice

©2013-2022 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 Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. 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

- Object Operations

 - Upload and Download

 - Data Management

 - Storage Class

- Logging and Monitoring

- 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

- Permission Management

- Data Processing

 - Function Service

- Data Security

- Tools

 - COSBrowser

 - COS Migration

 - COSCMD

 - COSFS

 - FTP Server

 - Hadoop

- SDKs

[FAQs](#)

[Android SDK](#)

[C SDK](#)

[.NET \(C#\) SDK](#)

[Go SDK](#)

[PHP SDK](#)

[iOS SDK](#)

[Python SDK](#)

[APIs](#)

FAQs

General

Last updated : 2022-04-28 14:34:01

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, 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 amount of data. There is no upper limit on bucket storage capacity and no partitioning is required. It supports HA deployment to ensure the eventual consistency of data but does not support features such as file locking. COS APIs support data access using HTTP or HTTPS protocol, and COS offers a variety of SDKs and tools that can be integrated into your business. Objects uploaded to COS can also be accessed or downloaded directly via a 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, file number, or storage capacity and needs to be managed and accessed via web APIs. COS offers various SDKs and tools for business integration, which can be used independently 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?

You can go to the object details page on the [COS console](#) to get the object address and signature link.

If your file is public-read:

- If you want to make it accessible to others all the time, we recommend directly using the object address.
- If you want to make it accessible to others for only a certain period of time, we recommend directly copying the signature link, which carries the signature parameter and is valid for 1 hour.

If your file is private:

- If you want to make it accessible to others all the time, we recommend changing its access permission to public-read/private-write and using the object address.
- If you want to make it accessible to others for only a certain period of time, we recommend directly copying the signature link, which carries the signature parameter and is valid for 1 hour.

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

Strictly speaking, the concepts of folders and directories do not apply to COS. However, taking into account the usage habits of different users, COS can display "folders" in the console just like the directory structure of a traditional file management system. 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. 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 in the COS console by deleting a single file, deleting files in batches, clearing incomplete multipart uploads, or emptying buckets.
- Deleting files using 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:
 - Downloading COS objects to your local file system or third-party servers using the [COSCMD tool](#).

- Performing 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.
- Using 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 data window. To view general data trends, go to the **Overview** page in the [COS console](#), and you can view data such as storage size, number of requests, and traffic of each storage class.

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

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?

The COS console has integrated Cloud Infinite features, which enable you to scale, crop, and add watermarks to your images. For more information, see [Enabling Image Processing](#).

Does COS support image compression?

COS is a distributed storage service for unstructured data and cannot support image compression on its own. For more information on image compression, see [Cloud Infinite](#).

Does COS support thumbnails?

COS is a distributed storage service for unstructured data and cannot support thumbnails on its own. For more information on thumbnails, see [Cloud Infinite](#).

Does COS support transcoding video files?

COS is a distributed storage service for unstructured data and cannot support video transcoding on its own. For information on how to transcode video files, see [Media Processing Service](#).

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 [File Decompression](#).

What are the specifications and limits of COS?

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

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](#).

Which version of COS should I use, an earlier version or the current version?

The implementation of earlier versions of COS is quite different than that of the current version. The current version has more features than earlier versions, and earlier versions are not updated with the latest features. **We recommend that you use the current version** for a better experience. If you are using an earlier version, [contact us](#) to activate the current version.

The current version comes with different APIs and SDK APIs than those in earlier versions. JSON APIs are used in earlier versions and XML APIs are used in the current version. JSON APIs have the same underlying architecture as XML APIs. Their data is interoperable and can intersect, but they are ultimately not compatible with each other and have different domain names.

How do I monitor error code information?

You can use [Cloud Monitoring](#) to get different types of HTTP error code messages. For more information, see [Monitoring and Alarms](#). For information on how to use cloud monitoring and obtain relevant data, see [Cloud Monitor Console Guide](#).

How do I calculate the availability of COS?

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

Tom uses COS to run his ecommerce business. Assume that his business incurred fees of 100 USD from November 1 to 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

2	15 min	November 15, 2018, 10:05 - 10:10	503	99	100
		November 15, 2018, 10:10 - 10:15	503	98	100
		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 the indemnification standard, Tom is eligible for indemnification 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 disable and/or stop being billed for COS services?

There is no one-click option for disabling COS. If there are long periods of time in which you do not use COS, you can save storage costs by transitioning your data to the ARCHIVE storage class. For more information on the transition operation, see [Setting Lifecycle](#).

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). There is no need to de-register your account, and if you use other Tencent Cloud products, avoid doing so as it will affect your other services.

Note :

To delete your COS data and incomplete multipart uploads, see [Deleting Objects](#), and [Deleting Incomplete Multipart Uploads](#).

Before the you complete the deletion process, note the following:

- Data, once deleted from COS, cannot be recovered, so make backups accordingly.
- Check your billing cycle to avoid overdue payments. 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.
- If your account has overdue payment and COS services are suspended, the free tier for which your account is eligible won't be available.

Billing

Last updated : 2022-12-07 17:44:35

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 public cloud prices for 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 after September 30, 2021).

2. Unit price of DEEP ARCHIVE read/write requests

- Product pricing: The public cloud prices for all 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 Mumbai, Seoul, Bangkok, Silicon Valley, Moscow, Jakarta, and São Paulo. The reduction has taken effect on July 1, 2022.

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 regions outside the Chinese mainland). For more information on regions, 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:

1. 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.
2. 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.
 3. 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](#).

Traffic

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

Public network downstream traffic is the 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 at a static website endpoint is public network downstream traffic. For billing details, see [Traffic Fees](#) and [Pricing | Cloud Object Storage](#).

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 public network downstream traffic generated after I enable CDN acceleration?

If you still use a COS origin domain name (in the format of `<bucketname-appid>.cos.<region>.myqcloud.com`) to access files in COS after enabling CDN acceleration, public network downstream traffic fees will be incurred. We recommend you use a CDN acceleration domain name instead, which will generate CDN origin-pull traffic only.

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 on the unit price, see [Traffic Fees](#).

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

CDN origin-pull traffic is the traffic generated by data transfer from COS to a CDN edge node. After CDN acceleration is enabled, traffic generated by browsing or downloading COS data in 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?

For data transfer between a COS bucket and a CVM instance in the same region, fees will be charged for requests but not the private network traffic. If they are in different regions, fees will be charged for both requests and traffic. For more information on how to identify private network access, see [Overview](#).

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 in the same region access each other over the private network by default, with no traffic fees incurred. For more information, see [Overview](#).

Bill


How do I view a 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 Billing Details](#). You can also 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](#).

L0: PDF Bills L1: Bill Summary L2: Bill by Instance L3: Bill Details

 L0 bills are in PDF format and can be used for payment requesting or archiving.

Period  to 

Account No

Download

What are billing by bucket and cost allocation by tag?

- Billing by bucket: It refers to using the bucket name as the **resource ID**, 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, billing by bucket has been enabled.

- ### (1) Bill by instance

Page 17 of 118

Bill by Instance

Bill Details

The current month's final bill for resource consumption will be generated on the 3rd day of the upcoming month. Price figures in Bill Details are accurate up to 8 decimal places. Expense figures in Bill by Instance are rounded off.

Cloud Object Storage
All Subproducts
Please choose one subproduct
All Projects

All Billing Modes
All transaction types
☐ Do not display \$0 transactions

Total Cost (Including Tax) **0.00 USD** = Total Amount After Discount (Excluding Tax) **0.00 USD** - Voucher

Instance ID	Instance Name	Transaction Time	Usage Start Time
kiwi-sz-125*****	kiwi-sz	2022-12-07 08:08:19	2022-12-06 00:00:00
kiwi-sz-125*****	kiwi-sz	2022-12-06 08:08:29	2022-12-05 00:00:00
kiwi-sz-125*****	kiwi-sz	2022-12-05 08:08:24	2022-12-04 00:00:00

[illegible]

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

- 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)
 - 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 about the billing cycle, please refer to [Billing cycle](#).

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	(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. (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.

Release Date	Release Note	Bill Description
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).

ProductName	BillingMode	Project	Region	Availability Zone	Inst	Instance	Subproduct	TransactionType	Transact	TransactionTime	Usage Start Time	Usage End Time	ComponentType	ComponentName
Cloud Object Storage	Pay-As-You-Go resourdefault		Western U.S. (Silicon Valley Zone 1)	100009C			cos standard	Monthly settlement	202209013	2022/9/1 8:14	2022/8/1 0:00	2022/8/31 23:59	standard storage usage	cos standard storage us
Cloud Object Storage	Pay-As-You-Go resourdefault		East China (Nanjing Zone 1)	100009C			cos standard	Monthly settlement	202209013	2022/9/1 8:14	2022/8/1 0:00	2022/8/31 23:59	standard storage usage	cos standard storage us
Cloud Object Storage	Pay-As-You-Go resourdefault		Southeast Asia (Singapore Zone 1)	100009C			cos standard	Monthly settlement	202209013	2022/9/1 8:14	2022/8/1 0:00	2022/8/31 23:59	standard storage usage	cos standard storage us
Cloud Object Storage	Pay-As-You-Go resourdefault		North China (Beijing Zone 2)	100009C			cos standard	Monthly settlement	202209013	2022/9/1 8:14	2022/8/1 0:00	2022/8/31 23:59	standard storage usage	cos standard storage us

- 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:

ProductName	BillingMode	ProjectN	Region	Availabi	Ind	N	Insta	Subprod	TransactionType	Transac	TransactionTime	Usage Start Time	Usage End Time	ComponentType	ComponentName
Cloud Object Storage	Pay-As-You-Go resoi default		North Chin	Beijing Zor	1000090t	-		cos standa	Daily settlement	20220902	2022/9/2 7:39	2022/9/1 0:00	2022/9/1 23:59	standard storage usage	cos standard storage
Cloud Object Storage	Pay-As-You-Go resoi default		East China	Nanjing Zc	1000090t	-		cos standa	Daily settlement	20220902	2022/9/2 7:39	2022/9/1 0:00	2022/9/1 23:59	standard storage usage	cos standard storage
Cloud Object Storage	Pay-As-You-Go resoi default		Southeast	Singapore	1000090t	-		cos standa	Daily settlement	20220902	2022/9/2 7:39	2022/9/1 0:00	2022/9/1 23:59	standard storage usage	cos standard storage
Cloud Object Storage	Pay-As-You-Go resoi default		Western U	Silicon Vall	1000090t	-		cos standa	Daily settlement	20220902	2022/9/2 7:39	2022/9/1 0:00	2022/9/1 23:59	standard storage usage	cos standard storage
Cloud Object Storage	Pay-As-You-Go resoi default		Southeast	Singapore	1000090t	-		cos standa	Daily settlement	20220903	2022/9/3 7:52	2022/9/2 0:00	2022/9/2 23:59	standard storage usage	cos standard storage
Cloud Object Storage	Pay-As-You-Go resoi default		North Chin	Beijing Zor	1000090t	-		cos standa	Daily settlement	20220903	2022/9/3 7:52	2022/9/2 0:00	2022/9/2 23:59	standard storage usage	cos standard storage
Cloud Object Storage	Pay-As-You-Go resoi default		East China	Nanjing Zc	1000090t	-		cos standa	Daily settlement	20220903	2022/9/3 7:52	2022/9/2 0:00	2022/9/2 23:59	standard storage usage	cos standard storage
Cloud Object Storage	Pay-As-You-Go resoi default		Western U	Silicon Vall	1000090t	-		cos standa	Daily settlement	20220903	2022/9/3 7:52	2022/9/2 0:00	2022/9/2 23:59	standard storage usage	cos standard storage

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.

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 payments?

After the COS service is suspended due to overdue payments, 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 ten 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.

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

COS comes with multiple billable items. Storage, requests, STANDARD_IA data retrievals, and ARCHIVE data retrievals are billed monthly, while DEEP ARCHIVE data retrievals, traffic, and management features are billed daily. For more information, see [Billing Overview](#).

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](#).

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

STANDARD_IA storage of data for less than 30 days will be calculated as 30 days, and storage for 45 days will be calculated as two cycles. For more information, see [Storage Usage Fees](#).

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

ARCHIVE storage of data for less than 90 days will be calculated as 90 days. For more information, see [Storage Usage Fees](#).

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

DEEP ARCHIVE storage of data for less than 180 days will be calculated as 180 days. If you successfully uploaded an object to DEEP ARCHIVE without enabling versioning, COS will delete the existing object (if any) that has the same name. In this case, storage usage fees will be incurred for the early deletion of the object. For more information, see [Storage Usage Fees](#).

What are data retrieval fees in COS?

Data retrieval fees are fees incurred by reading or downloading **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 STANDARD_IA data?

Copying 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](#).

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](#).

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 monthly fees incurred in the last month**. COS storage usage fees and request fees are billed monthly, that is, a bill generated in the current month is for usage in the last month. You can go to the [Transaction Details](#) page and click **Details** on the right of a billable item to view fee deduction details.

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](#).

Bucket Configuration Lifecycle

Last updated : 2022-10-08 17:48:37

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

Batch Operations

Last updated : 2022-10-08 17:43:34

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 : 2021-07-05 10:56:15

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 : 2021-07-05 10:56:15

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 : 2021-07-12 11:07:06

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 : 2022-04-28 14:34:01

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

Etag

Max-age *

300

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 -Lvo /dev/null "<object address="">" -H "origin:<domain name="">"`, for example, `curl -Lvo /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 : 2022-05-06 14:35:44

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. Currently, manual trigger is not supported. 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](#).

Object Operations

Upload and Download

Last updated : 2022-06-20 14:38:12

Does COS limit the upload and download bandwidth?

Yes. The default bandwidth threshold for a single bucket residing in a public region in the Chinese mainland is 15 Gbps, and 10 Gbps for those in other regions. If this threshold is exceeded, bandwidth throttling will be triggered for requests. To raise the threshold, [contact us](#).

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 `Content-Disposition` parameter cannot be `attachment`. If the browser supports the current file format, it will directly open the file instead of downloading it.

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 the `Content-Disposition` in the custom object headers to `attachment`. For detailed directions, see [Custom Headers](#).

You can also set the value of the request parameter `response-content-disposition` in the GET Object API to `attachment` so that your browser can pop up a window for the file to be downloaded. 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?

The access endpoints of COS use intelligent DNS resolution. For COS access via the Internet (including different ISPs), we will detect and select the optimal linkage for you to access COS. If you have deployed a service in Tencent Cloud to access COS, access within the same region will be automatically directed to a private network address. Cross-region access is not supported in a private network and the COS endpoint is resolved to a public network address by default.

How to determine access over a private network

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

to save on costs.

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 about 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, [contact us](#).

How do I download a folder?

You can log in to [COSBrowser](#), select the folder to be downloaded, and click **Download** to download the folder or files in batches. You can also download a 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 [A 403 Status Code is Returned When You Access COS](#).

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](#).
- APIs/SDKs: COS allows you to operate on multiple files by repeatedly calling an API or SDK. For more information, see [APIs for object uploads](#) and [SDK Overview](#).
- Tools: Use [COSCMD](#) and [COSBrowser](#) 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 in the console and "Upload failed due to network issues" 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?

You can set your bucket permission to private read/write. For more information, see [Setting Access Permission](#). You can also configure a hotlink protection allowlist on your bucket to block any access from endpoints that are 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 : 2022-10-11 12:08:12

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 : 2022-07-06 12:10:21

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). 0.025 USD is charged per 10,000 monitored objects per month.

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	$655.36 \times 0.024 + 368.64 \times 0.018 = 22.36$	$1024 \times 0.024 = 24.58$

Month	Object Monitoring Fees (USD)	INTELLIGENT TIERING Storage Usage Fees (USD)	STANDARD Storage Usage Fees (USD)
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 : 2022-04-28 14:34:01

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.cos.ap-chengdu.myqcloud.com DELETEObject 110.110.110.110 AKIDSuCmiBvppcdxShtPrCjhEUPFpzSzmXEEhG2bFVgd7-J6AsmEPu8NYM0hgX3HLEhx - 0 0 / DELETE tencentcloud-cos-console 200 - - 746 146 USER - 100009682373 - 100009682373:100009682373 NWU0MTU1NzRfNWNiMjU4NjRfM2JkMV8yNGFiNGEw - - - DELETE /filepath HTTP/1.1
```

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.

Domain Names and CDN

Static Website

Last updated : 2022-11-11 14:38:32

How do I display an object when I access it without downloading it?

Bind your own domain name and enable the static website hosting feature as instructed in [Setting Up a Static Website](#). If you don't have a domain name yet, you can [register](#) one through Tencent Cloud.

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 should I do if images cannot be displayed even after I enable the static website hosting feature?

1. If you want the objects (images) to be displayed directly when you access COS, you need to enable the static website hosting feature instead of setting `Content-Disposition:inline` in the object headers.
2. Check browser or CDN for cached data. You can use the curl and wget commands to avoid browser caching. The cached CDN URL can be purged in the [CDN console](#).

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 images cannot be displayed even after I enable the static website hosting feature?

Check browser or CDN for cached data. You can use the curl and wget commands to avoid browser caching. The cached CDN URL can be purged in the [CDN console](#).

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:
4. 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)

5. 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)

6. 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 : 2022-06-28 14:55:40

Does COS have a private network domain name?

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. You can use this domain name to access COS over a private network (COS will be resolved to a private IP). In this case, traffic generated is private network traffic and will not incur traffic fees.

An example domain name is `examplebucket-1250000000.cos.ap-guangzhou.myqcloud.com`. For more information on domains, see [Regions and Access Endpoints](#).

Global Acceleration Endpoint

Last updated : 2022-10-08 17:57:36

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

No.	API	No.	API
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 Tibet to the bucket in Beijing, the Tibet 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 Tibet to the bucket in Chengdu, the Tibet 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 : 2022-11-11 14:38:32

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 : 2022-11-30 15:55:43

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 : 2022-10-09 12:34:53

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, default CDN acceleration 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. It's important to distinguish it from the default CDN acceleration domain name.
- **Default accelerated domain name:** a domain name passing through CDN acceleration nodes and is generated by the system. You can choose to enable/disable it.
- **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.

Permission Management

Last updated : 2022-03-01 18:10:59

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, please 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, please see [Cases of Permission Setting](#).

What do I do if COS returns a 403 error?

You can use the [COS debugger](#) that the COS team developed for developers to configure Hosts to capture COS API requests and intelligently analyze the error in the requests. You can also troubleshoot the error according to the following process:

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 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](#).

You can also use the [COS debugger](#) that the COS team developed for developers to configure Hosts to capture COS API requests and intelligently analyze the error in the requests. The tool will provide the error details and handling suggestions.

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 incorrect permission error is reported when I am creating a bucket?

This is because your bucket is set to public read/private write, or public read/write. These permissions occupy the number of ACLs your root account is allowed to use. This error is reported because the root account has used up the

number of ACLs allowed, which cannot be adjusted.

The following two solutions are provided for your reference:

Solution 1: Set the access permission of the bucket to private read/write. For more information, please 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](#).

You can also use the [COS debugger](#) that the COS team developed for developers to configure Hosts to capture COS API requests and intelligently analyze the error in the requests. The tool will provide the error details and handling suggestions.

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, please see [Granting Sub-account Under One Root Account Permission to Manipulate Buckets 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

☐ Action Name

Description

☐ AbortMultipartU...

☐ AppendObject

☐ CompleteMultip...

☐ CreateJob

Create a COS Bat...

☐ DeleteBucket

☐ 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-1251*****

Operation *

Action Name	Operation
<div>GetBucket ▾</div>	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 access paths for the sub-account. For more information, see [Accessing Bucket List Using a Sub-account](#).

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, please see [CDN Origin-Pull Authentication](#).

Data Processing Function Service

Last updated : 2021-07-27 10:08:18

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 : 2021-10-08 15:23:43

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, please 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, please see [Versioning Configuration](#).
2. Using [cross-bucket replication](#) to achieve remote disaster recovery. For more information, please see [Cross-Bucket Replication Configuration](#).

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, please 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, please 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, please 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, please 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, please see [Common Request Headers](#).

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%.

Tools

COSBrowser

Last updated : 2022-06-28 15:45:00

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

SecretKey

Buckets/Access Path ⓘ

Region

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?

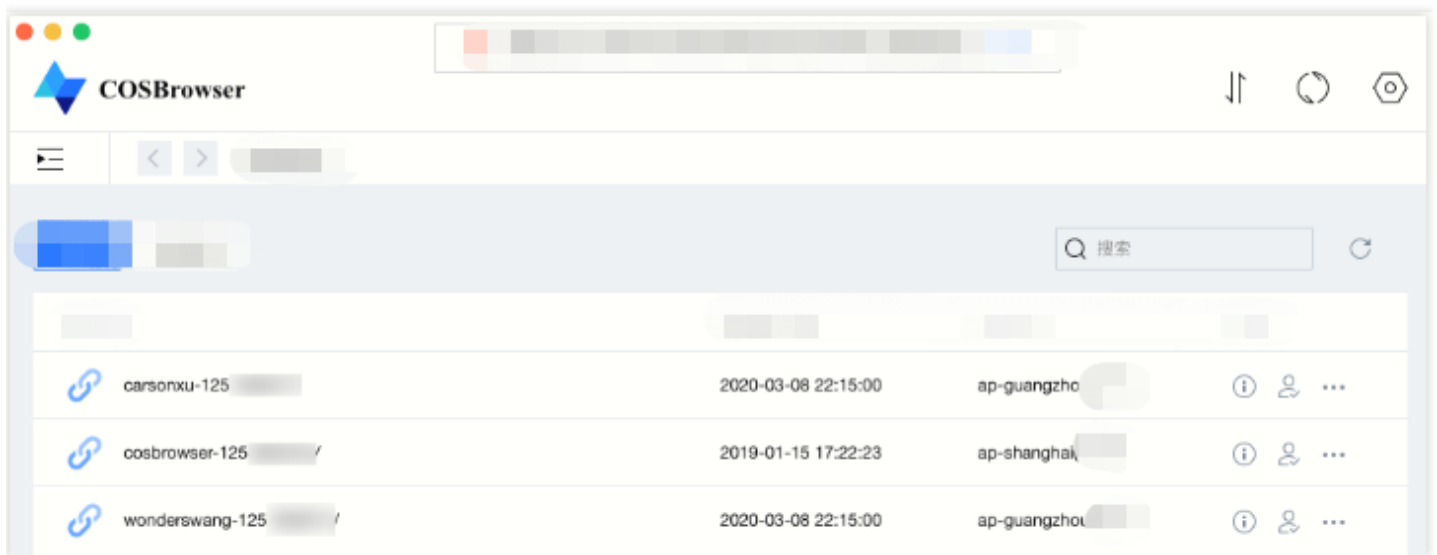
You can copy a file link as follows:

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

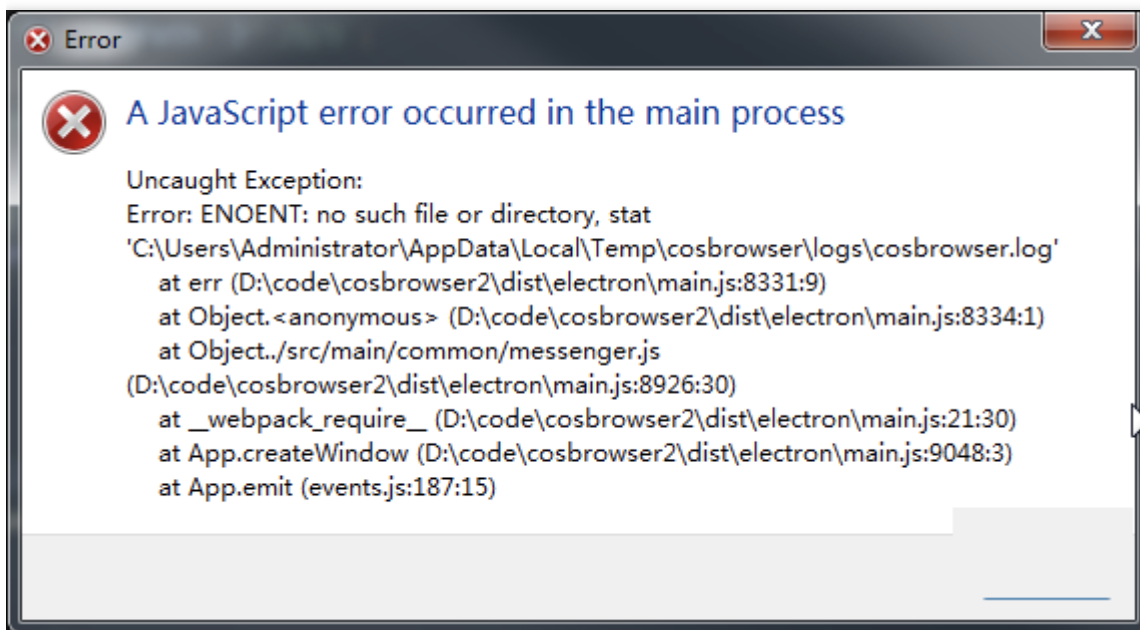
The above issue may occur when the two files `com.tencent.cosbrowser` and `com.tencent.cosbrowser.ShipIt` under the `/Users/username/Library/Caches/` directory belong to a root account and a user account, respectively. The resulting permission issue can cause update failures.

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 : 2022-06-09 16:12:12

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 `ddb` directory. Before the migration tool runs, the `ddb` directory will be scanned first, and the files recorded in it will not be uploaded again. For more information, please 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. You are advised to 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' name='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.

If other problems occur, try to run the migration tool again. If the problem persists, please compress the configuration (with the key hidden) as well as the log directory and [contact us](#).

COSCMD

Last updated : 2022-06-20 12:27:21

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.

COSFS

Last updated : 2022-11-30 14:59:53

FAQs About Features

How do I mount a bucket with a temporary key?

You can mount a bucket using a temporary key (STS) by following the steps below:

Step 1. Create a temporary key configuration file, such as `/tmp/passwd-sts`, for the COSFS command option `"-opasswd-file=[path]"` to specify the key configuration file. For more information on temporary keys, see [Generating and Using Temporary Keys](#). Below is an example of a temporary key configuration file:

```
COSAccessKeyId=AKIDXXXXXXXXXXXXXXXXXXXXXXXXXXXXX # Below are the Id, Key, and T
oken fields of the temporary key
COSSecretKey=GYYYYYYYYYYYYYYYYYYYYYYYYYYYYYYY
COSAccessToken=109dbb14ca0c30ef4b7e2fc9612f26788cadbfac3
COSAccessTokenExpire=2017-08-29T20:30:00 # Expiration time of the temporary token
in GMT time, which should be in the same format as shown here
```

COSFS will determine whether it needs to reload the configuration from the key file based on the time configured in `COSAccessTokenExpire`.

Note :

To prevent key leakage, you should set the permission of the key file to 600 in COSFS by running the following command:

```
chmod 600 /tmp/passwd-sts
```

Step 2: Run the COSFS command. Specify the role as sts using the command option `-ocam_role=[role]` , and specify the key file's path using `-opasswd_file=[path]` , as shown below:

```
cosfs examplebucket-1250000000 /mnt/cosfs -ourl=http://cos.ap-guangzhou.myqcloud.
com -odbglevel=info -oallow_other -ocam_role=sts -opasswd_file=/tmp/passwd-sts
```

How do I check the mounting parameter options and version number provided by COSFS?

You can check the parameter options provided by COSFS using the `cosfs --help` command, and check the COSFS version number using `cosfs --version`.

How do I view the logs generated by COSFS?

In CentOS, COSFS-generated logs are stored in `/var/log/messages`; in Ubuntu, these logs are stored in `/var/log/syslog`. If you have any problems during operation, send us the log for the corresponding time period.

How do I mount a directory in a bucket?

When running the mounting command, you can specify a directory under in bucket as shown below:

```
cosfs examplebucket-1250000000:/my-dir /mnt/cosfs -ourl=http://cos.ap-guangzhou.myqcloud.com -odbglevel=info
```

Note :

my-dir must begin with `/`.

In versions below v1.0.5, the mounting command is:

```
cosfs 1250000000:examplebucket:/my-dir /mnt/cosfs -ourl=http://cos.ap-guangzhou.myqcloud.com -odbglevel=info
```

How does a non-root user mount COSFS?

If you are not a root user, you are recommended to create a `.passwd-cosfs` file in the `Home` directory, set the permission to 600, and then mount COSFS using the mounting command. In addition, you can also specify the key file path using the `-opasswd_file=path` option and set the permission to 600.

Does COSFS support mounting over HTTPS?

Yes. Mounting methods over HTTP and HTTPS are as shown below:

```
-ourl=http://cos.ap-guangzhou.myqcloud.com  
-ourl=https://cos.ap-guangzhou.myqcloud.com
```

If the version of the NSS library that libcurl depends on is v3.12.3 or higher (check the NSS version using `curl -V`), you can mount the bucket over HTTPS by running the following command:

```
echo "export NSS_STRICT_NOFORK=DISABLED" >> ~/.bashrc  
source ~/.bashrc
```

How do I set auto-mounting at startup for COSFS?

You need to install the FUSE package first:

```
# CentOS
#sudo yum install -y fuse
# Ubuntu
#sudo apt-get install fuse
```

Add the following to the `/etc/fstab` file. The `_netdev` option specifies that the command can be executed only after the network is ready:

```
cosfs#examplebucket-1250000000 /mnt/cosfs fuse _netdev,allow_other,url=http://co
s.ap-guangzhou.myqcloud.com,dbglevel=info
```

How do I set the user and user group of files or directories under a mount point?

In certain scenarios (such as NGINX server), you need to set the user and user group of files or directories under a mount point, such as user `www` (uid = 1002, gid =1002). In this case, you need to add the following mount parameters:

```
-o uid=1002 -o gid=1002
```

How do I mount multiple buckets?

If you need to mount multiple buckets simultaneously, write a line for each bucket to be mounted in the `/etc/passwd-cosfs` configuration file, with the content of each line being in the same format as that for a single bucket, such as:

```
echo examplebucket-1250000000:AKIDXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX:GYYYYYYYYYYYYYY
YYYYYYYYYYYYYYYYY >> /etc/passwd-cosfs
echo log-1250000000:AKIDXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX:GYYYYYYYYYYYYYYYYYY
YYYYYY >> /etc/passwd-cosfs
```

How do I make a mounted directory accessible to other accounts on the server?

To make the mounted directory accessible to other accounts, specify `-o allow_other` during mounting.

Are there any restrictions on filenames in a COSFS-mounted directory?

In a COSFS-mounted directory, you can create a file whose name does not contain `/`. On a Unix-like system, the `/` character is a directory separator, so `/` is not allowed in filenames in a COSFS-mounted directory. In addition,

when creating a file whose name contains special characters, make sure that such characters are not used by shell; otherwise, file creation may fail.

How does COSFS determine whether a file exists?

In the internal logic of COSFS, a HEAD request is used to determine whether the parent directory and the file exist.

How can I view the storage usage with COSFS?

COSFS does not support viewing the storage usage. If you need statistics on bucket usage and your data volume is small, log in to the COS console to view the usage. If the data volume is high, you can leverage the [Inventory](#) feature.

How do I check which processes have accessed the mounted directory?

Run the following command to check which processes have accessed the mounted directory, such as

```
/mnt/cosfs .
```

```
ls -l /mnt/cosfs
```

FAQs About Troubleshooting

What should I do if the error message "unable to access MOUNTPOINT /path/to/mountpoint: Transport endpoint is not connected" is displayed and COSFS becomes inaccessible?

You can check whether the COSFS process exists using the `ps ax|grep cosfs` command. If the COSFS process is down due to faulty operations, run the following command to remount it:

```
umount -l /path/to/mnt_dir  
cosfs examplebucket-1250000000:/my-dir /mnt/cosfs -ourl=http://cos.ap-guangzhou.myqcloud.com -odbglevel=info
```

If the COSFS process error is not caused by faulty operations, you can check whether the fuse version of the server is below v2.9.4. The libfuse on versions below v2.9.4 can cause exceptional exit of the COSFS process. In this case, it is recommended to update the fuse version or install the latest version of COSFS as described in [COSFS](#).

What should I do if the Content-Type of a file uploaded via COSFS is changed to "application/octet-stream"?

COSFS automatically sets the Content-Type of a file uploaded to COS by comparing the file's suffix against `/etc/mime.types`. When a Content-Type problem occurs, check whether the configuration file exists on the system. For Ubuntu, you can add it using `sudo apt-get install mime-support`. For CentOS, you can add it using

`sudo yum install mailcap` . You can also create this file manually, where one line is added for each file format, as shown below:

```
image/jpeg jpg jpeg
image/jpm jpm jpgm
image/jpx jpx jpf
```

What should I do if "Bucket not exist" is displayed during mounting?

Check the parameter `-ourl` to ensure that the bucket part is not included in the URL. The correct format is:

```
-ourl=http://cos.ap-guangzhou.myqcloud.com
```

Why can't I write to files that I could write to previously?

Because of the adjustment to the COS authentication policy, using the COSFS tool below v1.0.0 will cause the policy check to fail. You can install the latest COSFS tool and mount it again.

What should I do if an error such as "Input/Output ERROR" occurs when using the COSFS tool?

Please follow the steps below to identify the cause of the error:

1. Check whether the server can access the COS domain name normally.
2. Check whether the account is configured correctly.
3. If you used the `cp` command containing the `-p` or `-a` parameter for copying, it is recommended that you remove the parameter and run the command again.

After confirming that the above configurations are correct, open the `/var/log/messages` log file on the server and locate the log entry for s3fs, which can help you identify the cause of error. If the error persists, [contact us](#).

I've set auto-mounting at startup for COSFS using /etc/fstab, but the error "wrong fs type, bad option, bad superblock on cosfs" occurs when I run "mount -a". Why?

This error happens generally because the fuse library is missing on your server. It is recommended to install the fuse library by running the following command:

- CentOS

```
sudo yum install fuse
```

- Ubuntu

```
sudo apt-get install fuse
```

Is it normal to see a lot of 404 error codes in the system log /var/log/messages?

In the internal logic of COSFS, a HEAD request is used to determine whether the parent directory and the file exist, and the 404 error does not necessarily mean that the program is running erroneously.

Why does the file that I see in COS have a size of 0?

Generally, when you write data to COSFS, a file with a size of 0 will be first created in COS, and the data will be written to the local cache file.

During the write process, the result of the mount point ls command shows the change in file size. When the file is closed, COSFS will upload the data written to the local cache file to COS. If the upload fails, you may only see a file with a size of 0, in which case you can try copying the failed file again.

Are the files in the COSFS cache directory the same as those in COS? Can I use them directly?

No. The files in the cache directory are used to accelerate reads and writes in COSFS and may include only a part of the files in COS.

I have copied a file to COSFS using the rsync command and the progress has reached 100%, but I only see a temporary file on the server. Why?

The rsync command creates a temporary file in the mounted directory. The 100% progress only means that the temporary file has been 100% written locally. After that, it will be uploaded to COS and then renamed and copied. Typically, it takes more time to copy data to a COSFS-mounted directory using the rsync command than using the cp command.

What should I do if COSFS uses up the disk capacity?

Uploads and downloads using COSFS involve disk file caching. When you upload or download a large file, if you don't specify the `-oensure_diskfree=[size]` parameter, the disk where the file is cached will be used up. You can use the `-oensure_diskfree` parameter to specify that when the remaining capacity of the disk is less than [size] MB, COSFS will minimize disk capacity usage (in MB). If you specify the `-ouse_cache=[path]` parameter, the cache file will be located in the path directory; otherwise, it will be in the /tmp directory.

For example, you can configure COSFS to reduce disk capacity usage when the remaining capacity is less than 10 GB by running the following command:

```
cosfs examplebucket-1250000000 /mnt/cosfs -ourl=http://cos.ap-guangzhou.myqcloud.com -odbglevel=info -oensure_diskfree=10240
```

What should I do if the error message "fuse: failed to open /dev/fuse: Operation not permitted" is displayed when I use docker to mount COSFS?

To start a docker image, you need to add the parameter `--privileged`.

Can I use a directory as a shared cache directory for multiple mount points?

It is not recommended that multiple mount points share a cache directory. The cache directory contains metadata used by COSFS, and sharing it may mess up the metadata.

What should I do if the error message "/bin/mount: unrecognized option --no-canonicalize" is displayed when I use COSFS for mounting?

Lower versions of the mounting tool do not support the `--no-canonicalize` option. Please update the tool. You are recommended to [download v2.17](#) and then mount it again. The installation command is as follows:

```
tar -jxvf util-linux-ng-2.17.tar.bz2
cd util-linux-ng-2.17
./configure --prefix=/usr/local/util-linux-ng
make && make install
mv /bin/mount{,.off}
ln -s /usr/local/util-linux-ng/bin/mount /bin
```

What should I do if mounting fails?

Step 1. Check whether the mounting command and the key configuration file are correct and whether the COS service is accessible based on the log files and error messages.

Step 2. Check whether the server time is correct by running the date command. If it is incorrect, correct it using the date command and then mount it again, such as `date -s '2014-12-25 12:34:56' .`

Is it normal that the time of a mounted directory changes to 1970-01-01 08:00 when I use the `ls -l --time-style=long-iso` command?

Yes. After you unmount the mount point, the time of the mounted directory will return back to the time before mounting.

Can a mounted directory be non-empty?

You can mount a non-empty directory using the `-ononempty` parameter, but you are not recommended to do so because a problem may occur when the mount point and the original directory have files with the same path.

Why does it take the `ls` command so long to return when I run it in a COSFS directory?

If there are a lot of files in a mounted directory, executing the `ls` command requires a HEAD operation on each file in the directory, so it takes a lot of time to read the directory system before the command returns.

Note :

You are recommended not to enable IO hung which may result in unnecessary restarts.

What should I do if the system log files generated take up a lot of storage space when I set the log level to info?

You can periodically clear the generated system log files or increase the log level. For example, you can use '-odbglevel=crit' for mounting.

What scenarios does COSFS apply to, and how about the read and write performance?

As disks are required for COSFS's read and write operations, COSFS is only suitable for scenarios where POSIX access syntax is required for COS access, such as machine learning algorithms of shared data sets reading shared data and simple log backups. COSFS adopts multithreaded upload and download for acceleration. Via a private network in the same region, it takes COSFS about 80s to sequentially read a 6 GB file and about 160s to sequentially write a 6 GB file. Generally, you can use SDK and multi-thread to achieve better performance.

Note :

The large number of system calls caused by file reading and writing, along with the large number of logs, can affect the COSFS read and write performance to some extent. If you have high performance requirements, you can specify `-odbglevel=warn` or higher log level.

What should I do if the system prompts that COSFS is not found after the COSFS RPM package is installed?

The installation path of COSFS is `/usr/local/bin`. If the system prompts that COSFS is not found, the possible cause is that the path is not specified in the `PATH` environment variable, and you need to add a line of configuration to `~/.bashrc` :

```
export PATH=/usr/local/bin:$PATH
```

Then you need to run the following command:

```
source ~/.bashrc
```

What should I do if the system reports "conflicts with file from package fuse-libs-*" when I install the COSFS RPM package?

Add the `--force` option when installing the COSFS RPM package:

```
rpm -ivh cosfs-1.0.19-centos7.0.x86_64.rpm --force
```

Why does the system indicate no permission when I separately mount a directory that has been granted read-only permission in COSFS?

COSFS requires the GetBucket permission on the root directory. Therefore, you need to add the GetBucket permission on the root directory and the read permission on the corresponding directory. In that way, you can list other directories but have no operation permission on them.

What should I do if COSFS has a high CPU utilization, sends a large number of HEAD and LIST requests to COS, and incurs a large amount of request fees during a certain period of time every day?

This is usually caused by scheduled disk scanning tasks on your server. A common disk scanning program on Linux is updatedb. You can add the COSFS-mounted directory to the `PRUNEPATHS` configuration item in updatedb's configuration file `/etc/updatedb.conf` to prevent disk scanning. In addition, you can use the Linux tool auditd to find programs that access the COSFS mount point.

Follow the steps below:

Step 1: Install auditd.

Ubuntu:

```
ap-get install auditd -y
```

CentOS:

```
yum install audit audit-libs
```

Step 2: Start the auditd service.

```
systemctl start auditd  
systemctl enable auditd
```

Step 3: Monitor the mounted directory.

Note :

`-w` specifies the COSFS-mounted directory, and `-k` indicates the key to be outputted to the audit logs.

```
auditctl -w /usr/local/service/mnt/ -k cosfs_mnt
```

Step 4: Based on the logs, determine the programs that access the COSFS mount point.

The audit logs are in the `/var/log/audit` directory and can be queried by running the following command:

```
ausearch -i | grep 'cosfs_mnt'
```

Step 5: Stop the auditd service.

You can run the following command to stop the auditd service:

```
/sbin/service auditd stop
```

Note :

If the program accessing the mount point is always running, the newly started auditd will not monitor the access behavior of the program. This is because only the first call in multiple calls from the program to the mounted directory will be recorded.

Why are the values of `Size` and `Available` are 256 TB after I run `df` ?

COS buckets offer unlimited storage capacity. The 256 TB displayed is only used as the output of `df` .

Why is the value of `Used` is 0 after I run `df` ?

COSFS does not occupy local storage. In order to be compatible with tools such as `df` , the values of `Size` , `Used` , and `Available` displayed in COSFS are not the actual values.

Why are the values of `Inode` , `IUsed`, and `IFree` are 0 after I run `df -i` ?

COSFS is not a disk-based file system and thus does not have `inode` .

FTP Server

Last updated : 2021-08-13 11:06:17

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 : 2021-12-16 12:59:28

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 with 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, please 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/${user.name}/.Trash` directory, but no actual deletion will occur, so the data will still remain in COS. You can use the `-skipTrash` parameter to skip the recycle bin feature and delete the data directly. To implement periodic data deletion like with the HDFS recycle bin, please configure a lifecycle rule for objects prefixed with `/user/${user.name}/.Trash/`. For the configuration guide, please 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

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`.

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, a high level of concurrency may trigger the COS frequency control policy, resulting in a 503 `Reduce your request rate` error. You can initiate retries for the failed requests by configuring the `fs.cosn.maxRetries` parameter, which defaults to a maximum of 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 a tag of 5.8.3 or above. For more information, please see [Hadoop-COS on GitHub](#).

Parts

How do I set a reasonable part size for a multipart upload through Hadoop-COS?

Hadoop-COS uploads large files to COS through 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, i.e. those larger than the part size (`fs.cosn.upload.part.size`), through multipart upload. You can see the file on COS only after all of its parts have been 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, this exception occurs when you have established too many TCP non-persistent 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. During the timeout period, there will be no ports available 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 TIME_WAIT status t
o form a new TCP connection
net.ipv4.tcp_tw_recycle = 1 #Enables quick repossession of a socket in the TIME-W
AIT status
net.ipv4.tcp_syncookies=1 #Enables SYN Cookies. The default value is 0. When the
SYN waiting queue overflows, cookies are enabled to prevent a small number of SYN
attacks.
net.ipv4.tcp_fin_timeout = 10 #Waiting time after the port is released.
net.ipv4.tcp_keepalive_time = 1200 #The time interval between which TCP sends Kee
pAlive messages. The default value is 2 hours. Change it to 20 minutes.
net.ipv4.ip_local_port_range = 1024 65000 #The range of ports for external connec
tions. The default value is 32768 to 61000. Change it to 1024 to 65000.
net.ipv4.tcp_max_tw_buckets = 10240 #The maximum number (default value: 180000) o
f sockets in the TIME_WAIT status. Exceeding this number will directly release al
l the new TIME_WAIT sockets. You may consider reducing this parameter for a small
er number of sockets in the TIME_WAIT status.
```

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.LinkedBlockingQueue.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>
```

SDKs

FAQs

Last updated : 2022-06-20 10:35:29

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 `QCloudCredentialFenceQueueDelegate` protocols during initialization. The `QCloudCredentialFenceQueue` 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 `/`.

Android SDK

Last updated : 2021-07-27 10:08:18

What do I do if the client network is normal, while the access to COS over HTTP is very slow, or the error message "Connection reset" is reported?

In some regions, carriers may hijack COS endpoints. Therefore, you are advised to access COS over HTTPS.

What do I do if `ETag` is not included in the `Complete Multipart Upload` request and an error message "400 Bad Request" is reported?

The possible cause is that the `ETag` header is filtered out by the network. After the parts are uploaded, the SDK fails to parse the `ETag` parameter and reports the error in the Complete Multipart Upload operation.

What do I do if `QCloudResultListener` or other callback functions did not work?

If you view logs to determine whether a callback function works, the possible cause is that the log level is set too high, or the desired log is filtered out by other filtering rules. You can adjust the filtering rule, or set breakpoints in the callback function to determine whether the callback function works properly.

What do I do if `NoClassDefFoundError` is reported when I call an API?

The SDK depends on the bolts and OkHttp common classes. If the methods in these classes cannot be found, you might have imported these two dependencies to your project, and the version number is too low. You are advised to use a version consistent with the SDK version or a later one.

What do I do if the SDK failed to obtain permissions to the phone?

To upload files to or download files from an external storage device, you must have network permission and read/write permissions on the device. Other permissions, such as location permission and device information permission, are not mandatory. If you have strict requirements on permissions, you can skip importing the MtaUtils package or upgrade the SDK to 5.5.8 or above.

What do I do if the `java.security.cert.CertPathValidatorException: Trust anchor for certification path not found` error is reported when HTTPS is used?

If you access COS via a proxy, check whether the proxy supports HTTPS. If not, please [contact us](#).

What do I do if the upload progress reaches 100%, while the `onFailed` method is called?

The 100% upload progress indicates only the SDK packet sending progress. The upload is successful only when the `onSuccess` method is called. If an exception occurs when the `Complete Multipart Upload` request is sent, the `onFailed` method will be called.

What do I do if an error, such as `400 Bad Request` and `409 Conflict`, occurs in a multipart upload?

Use the advanced API `TransferManager` provided by the SDK for upload/download if possible. Encapsulating the multipart upload API may easily cause an error.

What do I do if a permission error is reported when I use `TransferManager` for upload/download?

The Head operation will be performed when `TransferManager` downloads an object. Therefore, the `HeadObject` and `GetObject` permissions should be granted for the download. For an upload operation, permissions on all simple upload and multipart upload APIs should be granted.

What do I do if an error, such as `lock timeout`, `no credential for sign`, or `expired signature`, is reported?

If you have implemented the `BasicLifecycleCredentialProvider#fetchNewCredentials()` method, please check whether the key is updated in time, or whether it is still valid. For a temporary key, the token should be carried.

What do I do if the `java.lang.RuntimeException: Can't create handler inside thread that has not called Looper.prepare()` error is reported?

If the error is reported when the `TransferManager#upload()` method is called in the master thread, this is a false positive reported by the MTA and can be ignored. You can also upgrade your SDK to 5.5.8 or above to solve this issue.

What do I do if an application error is reported when I directly operate the UI in a callback?

The SDK callback thread is not necessarily the master thread. Please do not operate the UI directly.

What do I do if `calculate md5 error` is reported during upload?

The possible cause is either that you have modified the file during the upload, changing the MD5 checksum, or the network is poor, causing a packet receive error on the server.

What do I do if `ServerError` is returned for a request?

If you access COS using a proxy, the possible cause is that the proxy returned the incorrect packet, causing the SDK to fail the parsing. You can capture the packet received by the client to verify the packet.

What do I do if the `403` permission error is reported when I call an API?

In general, a permission error is irrelevant to SDK. You can check your permissions or [contact us](#).

Does Android SDK support checkpoint restart?

Advanced APIs of the Android SDK of COS support checkpoint restart. To implement checkpoint restart, refer to the descriptions of the advanced APIs in [Uploading and Copying Objects](#).

C SDK

Last updated : 2022-04-28 14:34:02

How do I implement checkpoint restart with C SDK?

You can use the [advanced upload API of C SDK](#) to implement the checkpoint restart feature. To use checkpoint restart, you need to set the upload control parameter to **COS_TRUE**, for example, `clt_params = cos_create_resumable_clt_params_content(p, 0, 1, COS_TRUE, NULL)`.

Why does the `HttpIoError` error occur when I use C SDK?

Error description: When you use the SDK, all APIs cannot be used or return `requestid`. By analyzing the captured packets, you find that no HTTP requests are sent successfully as shown in the following logs:

```
transport failure curl code:1 error:Unsupported protocol
status->code: -996
status->error_code: HttpIoError
status->error_msg: Unsupported protocol
status->req_id:
```

This error occurs because the HTTPS protocol is used, but the libcurl library doesn't support HTTPS. Therefore, the OpenSSL library is not used or the versions mismatch during libcurl compilation.

Solution: Check the running environment and reinstall the libcurl library (if you install it by compiling the source code, enable SSL) or update the OpenSSL library.

.NET (C#) SDK

Last updated : 2021-07-05 11:06:57

Does the .NET SDK support downloading an object to memory?

Yes.

Can I use the .NET SDK to obtain the file stream directly?

Currently, you can download an object to a local file system or the memory. If you need to obtain the file stream, you can download the object to local and open a local file stream.

Go SDK

Last updated : 2022-10-08 12:43:39

How do I read all attributes of a file with Go SDK?

You can use the [GET Object](#) API to obtain file attributes from `Response.Header` .

Does Go SDK support file stream upload?

Go SDK allows you to upload file or byte streams. For more information, see [Uploading Objects](#).

PHP SDK

Last updated : 2021-06-09 14:27:13

What do I do if the `Call to undefined method`

`GuzzleHttp\Utils::chooseHandler()` error is thrown when the PHP SDK is running?

The PHP SDK depends on Guzzle. You are advised to install the SDK using Composer.

- If you install the SDK using Composer, run the `php composer.phar install` command to install the SDK and the dependencies.
- If you install the SDK using the source code, run the `composer install` command to install the SDK and the dependencies. For more information, please see [PHP SDK - Download and Installation](#).

What do I do if the `cURL error 60: See http://curl.haxx.se/libcurl/c/libcurl-errors.html` error is reported when I use the PHP SDK to upload files?

If something is wrong with your certificate in the PHP environment, errors similar to `cURL error 60: See http://curl.haxx.se/libcurl/c/libcurl-errors.html` will be reported. You can solve the problem as follows:

1. Download the `cacert.pem` certificate file at `https://curl.haxx.se/ca/cacert.pem` and save it to the PHP installation path.
2. Edit the `php.ini` file, delete the semicolon (;) before the `curl.cainfo` configuration item, and set its value to the absolute path of the `cacert.pem` certificate file.
3. Restart PHP-dependent services.

iOS SDK

Last updated : 2021-09-23 12:22:40

What do I do if the error `[__NSCFConstantString matchesRegularExpression]: unrecognized selector sent to instance xxx` is thrown after I manually integrate the SDK and set `regionName` for the `QCloudCOSXMLEndPoint` instance?

Cause: `matchesRegularExpression` is a method in the `NSString` category. Objective-C does not define linker symbols for each method, but for each class. The SDK is a static library. If a category has been defined for an existing class in the static library, the system assumes that the class already exists and will not integrate the code of the class and category. As a result, methods in the category will be missing in the executable code.

Solution:

1. Select the target, click **Build Settings > All > Other Linker Flags**, and add the `-Objc` and `-all_load` parameters.
2. `-Objc` loads the classes and categories in the static library to the executable file. If there are only categories but no classes in the library, `-Objc` will not take effect. In this case, `-all_load` is needed to load all object files to the executable file.

What do I do if the error `Default OCR configuration is not set/Key is not set to 'xxx' in the OCR configuration. Please call this method after it is configured` is thrown after I integrated the SDK and sent a request?

Cause: All SDK requests depend on `QCloudCOSXMLService` and `QCloudCOSTransferMangerService` (advanced upload APIs depend on this instance). If the corresponding service is not registered before the request is sent, this error will be thrown.

Solution: Use the following code to register the service instance required by the request. Ensure that the instance exists before you send a request.

```
QCloudServiceConfiguration *configuration = [QCloudServiceConfiguration new];
configuration.appID = "AppId";
configuration.signatureProvider = self;
QCloudCOSXMLEndPoint *endpoint = [[QCloudCOSXMLEndPoint alloc] init];
endpoint.regionName = @"region";
endpoint.useHTTPS = YES;
configuration.endpoint = endpoint;
[QCloudCOSXMLService registerDefaultCOSXMLWithConfiguration:configuration];
[QCloudCOSTransferMangerService registerDefaultCOSTransferMangerWithConfiguration:configuration];
```

What do I do if the error `Default COSXMLService already exists. To use new configurations, re-register using registerCOSXMLWithConfiguration:withKey:` is thrown after I integrated the SDK?

Cause: In the SDK, different `COSXMLService` instances correspond to different configurations (for the configurations, please see relevant attributes). For example, setting `regionName` to `guangzhou` and `beijing` indicates two different configurations, where two services need to be registered. If you already used `ap-guangzhou` to register a service and then set `region` to `ap-guangzhou` to register again, this error will occur.

Solution:

1. The default service uses `registerDefaultCOSXMLWithConfiguration:` and the key does not need to be specified.
2. Use the following code to register a new service:

```
// Determine whether the key to register exists.
if (![QCloudCOSXMLService hasServiceForKey:@"Key to register"]){
    // If the key does not exist, register a new service.
    [QCloudCOSXMLService registerCOSXMLWithConfiguration:configuration withKey:@"Key to register"]
}
```

What do I do if the `-(void) signatureWithFields: (QCloudSignatureFields *) fields request:urlRequest:complete:` proxy is not called after I run the SDK?

Cause: This proxy method obtains the key/signature. To give full play to the key validity (calling the signature earlier will cause it to expire sooner), the SDK calls the key only before the request is sent (`[task resume]`).

Solution:

1. Ensure that the class where the proxy method belongs is not terminated before the request is sent. You are advised to implement the proxy method in a singleton class.
2. Check whether `signatureProvider` is configured (for example, configured as `configuration.signatureProvider = self`, in which `self` is the class where the proxy method belongs) after the `QCloudServiceConfiguration` instance is created.
3. After checking the two points above, check whether the request is sent.

Can the SDK cache or reuse the key? How can I request a new key when the key expires?

The SDK provides `QCloudCredentialFenceQueue` to cache or reuse the key. For more information, please see [Getting Started](#).

What can I do if the SDK advanced API `QCloudCOSXMLUploadObjectRequest` throws the error `Body of this type is not supported`. Supported types are `NSData`, `QCloudFileOffsetBody`, and `NSURL` ?

Currently, the SDK supports the following three types of bodies:

1. `NSURL`: local path of the file. The URL is initialized through `[NSURL fileURLWithPath:@"Local path of the object"]`
2. `NSData`: binary data
3. `QCloudFileOffsetBody`: multipart body. Since it is the body type for internal use by SDK advanced APIs, you can ignore this type in most cases.

What do I do if I use the SDK advanced API `QCloudCOSXMLUploadObjectRequest` to upload videos or files from the system photo library, but checkpoint restart fails, throwing the error `The specified Content-Length is zero` ?

The SDK supports checkpoint restart only for files in the sandbox. To use the checkpoint restart feature, move your files to the sandbox first.

What do I do if the size of the file uploaded through an API is different from that of the local file after I integrated the SDK?

Ensure that the local file is not modified after you set the body. For example, if the upload API is called when the file is still being compressed or file write is not completed yet, the SDK will upload the file (using multipart upload) according to the size at the time of upload, causing a file size inconsistency between the COS file and the local file.

What do I do if I integrate the SDK and call the upload API, but the error “URL in the body is not a local URL” is reported?

Solution:

Ensure that the URL starts with “file://”. You can initialize in either of the following two ways:

1. `[NSURL URLWithString:@"file:///var/mobile/Containers/Data/Application/DBPF7490-D5U8-4ABF-A0AF-CC49D6A60AEB/Documents/exampleobject"]`
2. `[NSURL fileURLWithPath:@"file:///var/mobile/Containers/Data/Application/DBPF7490-D5U8-4ABF-A0AF-CC49D6A60AEB/Documents/exampleobject"]`

What do I do if I integrate the SDK and call the upload API, but the size of the successfully uploaded file is 0?

Solution:

1. If the path of the system photo library is used, check whether you have read permission on the file. An example is that the path `file:///var/mobile/Media/DCIM/101APPLE/` cannot be accessed directly. Photos in this path can only be obtained using the request method in the Photos framework.

2. To make your app compatible with iOS 11, call the `[[PHImageManager defaultManager] requestPlayerItemForVideo:asset options:option resultHandler:^(AVPlayerItem *playerItem, NSDictionary *info) { // Move the file to the sandbox in time or save the playerItem. }];` method to save the obtained `playerItem`, or move the desired files to the App Sandbox during the callback. The cause of this error is that for iOS11, when `playerItem` is released, read permission on the file specified in `playerItem` will expire, making the size of the uploaded file 0.
3. If the uploaded file is in the App Sandbox, check whether the file is stored in the `tmp` directory (e.g., `/var/mobile/Containers/Data/Application/0BFBB3FE-0FD0-46CB-ADDE-DDE08F6D62C3/tmp/`). Files in `tmp` will be cleared by the system any time. You can move your files to a securer directory in case they are cleared during the upload. For more information about sandbox, please see [File System Basics](#).

What do I do if I use an SDK advanced API for upload, but the error `The MD5 checksum is inconsistent with the local file. Please check whether the file is modified during the upload. During multipart uploads, the MD5 checksum of each uploaded part will be verified against the local file. Any inconsistency will cause an error` is reported?

Cause: If you use the SDK to upload a file greater than 1 MB, the file will be divided into multiple 1-MB parts and uploaded using multipart upload. After all parts are uploaded, the value of the backend-returned `ETag` will be compared with that of local parts. If there is any inconsistency, this error will be thrown.

Solution: Check whether the file is modified during the upload.

How can I set the request timeout in the SDK?

Solution: SDK 5.7.0 and later versions support customizing request timeout as follows:

1. Initialize `QCloudServiceConfiguration *config = [QCloudServiceConfiguration new]`.
2. Set the `timeoutInterval` attribute of `config`, for example, `.timeoutInterval = 30;`.

What should I do if the iOS SDK crashes?

Check whether you are using the latest version. If not, you are advised to update to the latest SDK version. For the latest iOS SDK, please see [qcloud-sdk-ios](#).

Python SDK

Last updated : 2022-11-11 14:38:32

What should I do if I cannot move objects after upgrading the Python SDK?

You can call the [PUT Object - Copy](#) and [DELETE Object](#) APIs to move objects. We recommend you check the data consistency before deleting an object as instructed in [MD5 Verification](#).

How do I get a temporary download URL with the Python SDK?

The Python SDK provides APIs for getting request signatures, pre-signed URLs, and pre-signed download URLs. The calling methods to get a pre-signed URL by using a permanent key or temporary key are the same, except that if you use a temporary key, you need to add `x-cos-security-token` to the header or `query_string`. For more information, see [Getting Pre-Signed URLs](#).

What should I do if the Python SDK reports an exception?

If an XML Python SDK operation is successful, a dict or None will be returned. If an SDK API for requesting the COS service fails, the system will report a `CosClientError` or `CosServiceError`.

- `CosServiceError`: Errors returned due to client requests not meeting requirements, such as accessing a non-existent or unauthorized object. For more information, see [Error Codes](#).
- `CosClientError`: Network exceptions, I/O errors during file reads/writes, parameter verification failures, etc.

APIs

Last updated : 2021-07-19 10:21:32

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.