Content Delivery Network

FAQ

Product Documentation

Tencent Cloud

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

FAQ

Features

Last updated：2023-10-11 15:14:27

Does Tencent Cloud CDN support global acceleration?

Yes. Tencent Cloud CDN supports acceleration in and outside the Chinese mainland, and global acceleration. Tencent Cloud CDN offers more than 800 nodes in more than 70 countries and regions outside the Chinese mainland to help accelerate your global business seamlessly. To implement acceleration outside the Chinese mainland or global acceleration, we recommend that you deploy your origin server outside the Chinese mainland to ensure smooth acceleration. If the origin server is deployed in the Chinese mainland but the acceleration region resides outside the Chinese mainland, or if the origin server is deployed outside the Chinese mainland but the acceleration region resides in the Chinese mainland, the acceleration effect of cross-border origin-pull may not meet your expectations.

After connecting to CDN, do changes need to be made on the origin server for the acceleration service to take effect?

No. To achieve a better acceleration result, however, you are recommended to assign static and dynamic files to different domain names and only accelerate static resources.

Does Tencent Cloud CDN support cross-region access?

Yes. If cross-region access is needed for your website, configure the `Access-Control-Allow-Origin` field on your website or configure cross-region headers for your domain name in the CDN console. For more information, see HTTP Response Header.

How do I use the CDN self-diagnosis tool?

CDN provides a self-diagnosis tool. If a URL cannot be accessed, you can use the tool to identify the cause of the failure by checking the DNS resolution configuration, cache nodes, and origin server network. The tool also offers solutions to help you troubleshoot the failure.

Does CDN support POST requests?

Yes.

Does CDN support the Cache-Control configuration of the origin server?

Yes. By default, CDN supports the origin server's Cache-Control configuration.

Does CDN support gzip compression?

Yes. To help you save traffic, CDN compresses files between 256 bytes and 2,048 KB with filename extensions of .js, .html, .css, .xml, .json, .shtml, and .htm into gzip files.

Can I customize the access port for CDN acceleration?

By default, access ports 80, 443, and 8080 are enabled for CDN acceleration. You can disable any of them as needed.

What is a CDN intermediate server?

A CDN intermediate server is a middle-layer origin-pull server located between the business server and the CDN node. The intermediate server converges the node's origin-pull requests to reduce the origin-pull pressure on your origin server.

How do I obtain the IP of a client that sends a request to the origin server and the protocol that is used by the client to send the request?

After a request goes through an edge node for acceleration, Tencent Cloud CDN adds the `X-Forwarded-For` and `X-Forwarded-Proto` headers to the request by default before the request is sent to the origin server. The `X-Forwarded-For` header indicates the IP of the client that sends the request. The `X-Forwarded-Proto` header indicates the protocol that is used by the client to send the request. You do not need to configure them.

How do I configure CDN sub-users?

Sub-users do not need to register for Tencent Cloud or activate the CDN service. They are added to the sub-user list by the creator. There are two types of sub-users:
1. Message recipients.
2. Console users. For more information on how to create and configure a sub-user, please see Console Permissions.

How do I configure an IP blocklist/allowlist in CDN?

CDN supports IP blocklist/allowlist configuration. You can create filtering policies for source IPs of user requests based on your business needs, helping prevent hotlinking and attacks from malicious IPs. For more information, please see IP Blocklist/Allowlist Configuration.

For more information on this configuration, please see IP Access Limit Configuration and Hotlink Protection Configuration.

Is there a size limit for a file uploaded to CDN?

Yes. The maximum size of a file that can be uploaded to CDN is 32 MB by default.

Does CDN support dynamic origin-pull configuration and origin-pull queuing?

If the primary origin server responds exceptionally, it can redirect requests to the configured backup origin server in sequence for request again.

Does CDN permanently block access to a blocked URL?

No, CDN does not permanently block access to a blocked URL.

Does CDN support WebSocket?

ECDN domain names support WebSocket. You can enable WebSocket in the advanced settings on the domain name management page.

Does CDN support acceleration by using protocols other than HTTP?

Yes, CDN not supports acceleration by using non-HTTP protocols, such as email protocols and FTP.

Billing

Last updated：2020-11-23 14:32:26

Can CDN be billed by the number of requests?

No. Currently, CDN does not support billing by the number of requests.

What will happen if my account falls into arrears?

For more information, please see the note on arrears in the billing description document.

If the origin server uses COS, will I be charged for traffic generated by origin-pull from CDN to COS?

No. The traffic generated by origin-pull from CDN to COS is billed by COS instead of CDN. For more information, please see COS as Origin Server.

When CDN is disabled (or deactivated), will it generate traffic and incur fees?

After the CDN domain name acceleration service is disabled, if the domain name is still configured with CNAME, a 404 status code will be returned for requests resolved to the node and a small amount of traffic will be consumed. The console will record this traffic data for your reference. Corresponding logs will also be generated. However, since your domain name has been disabled, you will not be billed for this traffic consumption and log packets. We recommend you modify the origin-pull resolution first before disabling the acceleration service.

How do I change the billing mode of CDN?

If you find the selected billing mode unsuitable for your actual business conditions (for more information on how to select the right billing mode for you, please see Choosing a Billing Plan), you can change it by following the steps below:

1. Log in to the CDN Console, access the service overview page, and click **Change** in the billing mode section on the right.

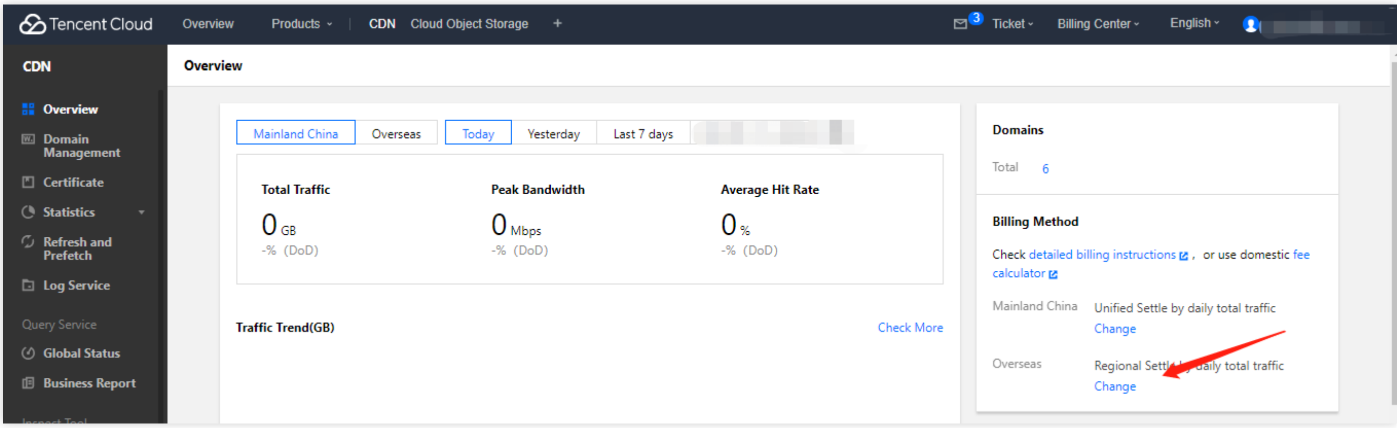

Billing Method

Check detailed billing instructions, or use domestic fee calculator

Mainland China	Unified Settle by daily total traffic Change
Overseas	Regional Settle by daily total traffic Change

2. Change the original billing mode **Bill by Traffic** to **Bill by Bandwidth**.

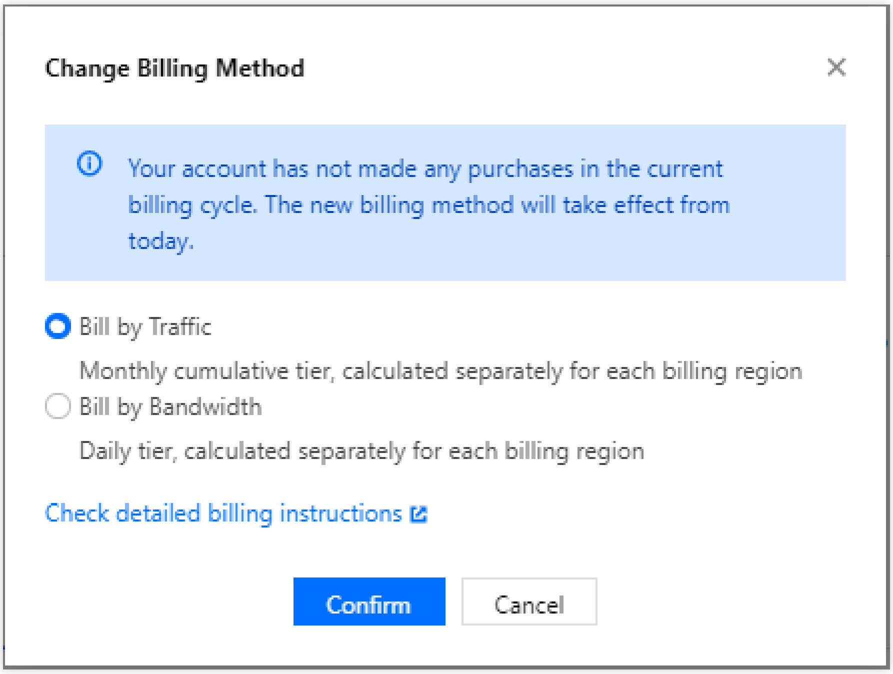

Change Billing Method

Your account has not made any purchases in the current billing cycle. The new billing method will take effect from today.

● Bill by Traffic

Monthly cumulative tier, calculated separately for each billing region

○ Bill by Bandwidth

Daily tier, calculated separately for each billing region

Check detailed billing instructions

3. After you change the billing mode to **Bill by Bandwidth**, the charge for total consumption generated on the current date will be billed and deducted on the next day:

If the origin server uses CVM, will I be charged for traffic generated by origin-pull from CDN to CVM?

No. CDN does not charge for this type of traffic.

What is the conversion rate from Gbps to Mbps in CDN billing?

In CDN billing, 1 Gbps = 1000 Mbps, 1 Mbps = 1000 Kbps, and 1 Kbps = 1000 bps.

Is only downstream traffic billable in CDN?

Yes. CDN only charges for downstream, not upstream traffic.

Is there a delay in using APIs to query data? How long is it?

There is a certain delay in using APIs to query data. Queries of real-time data such as access data and billing data have a delay of around 5–10 minutes, while queries of analytical data such as rankings will have delays of approximately half an hour. The data is calibrated on the backend at around 3 am Beijing Time.

FAQs about Domain Name Connection

Last updated：2023-03-10 15:02:35

How do I connect a domain name?

You can connect a domain name in the Content Delivery Network (CDN) console. For more information, see Adding Domain Names.

Are there any requirements for connecting a domain name to CDN?

1. The domain name cannot exceed 81 characters in length.
2. If the domain name requires acceleration in the Chinese mainland or global acceleration, you must obtain an ICP filing for the domain name from the Ministry of Industry and Information Technology (MIIT). If the domain name requires acceleration outside the Chinese mainland, you do not need to obtain an ICP filing for the domain name.
3. It takes 1 to 2 hours to synchronize ICP filing information. Re-add the domain name 1 or 2 hours after the ICP filing is complete.
4. The domain name can contain underscores (_) and Punycode-converted Chinese characters. You must obtain ICP filings for Chinese domain names before you convert the Chinese characters in the domain names to Punycode.
5. You can connect wildcard domain names in various formats, such as `*.example.com` and `*.a.example.com` . After you connect a wildcard domain name, its subdomain names or second-level wildcard domain names cannot be connected under another account. For example, if you connect the `*.example.com` wildcard domain name, the access traffic to the `a.example.com` domain name is accelerated based on the settings that are configured for the `*.example.com` wildcard domain name, whereas the access traffic to the `example.com` domain name is not accelerated.
6. You can connect multiple levels of nested domain names, such as `*.example.com` , `*.path.example.com` , and `a.path.example.com` , at the same time under an account. In this case, the domain name settings are applied and traffic statistics are calculated based on the priorities of the domain names. A more accurate match between the accessed domain name and a connected domain name indicates a higher priority for the connected domain name. For example, the access traffic to `a.path.example.com` is accelerated based on the settings of `a.path.example.com` , the access traffic to `b.path.example.com` is accelerated based on the settings of `*.path.example.com` , and the access traffic to `c.example.com` is accelerated based on the settings of `*.example.com` . The same analogy applies to traffic statistics.
7. If the subdomain names of a wildcard domain name are connected under other accounts, the wildcard domain name can be connected only after the subdomain names are disconnected under the accounts. For example, if the `a.example.com` subdomain name of the `*.example.com` wildcard domain name is connected under Account A, you must delete the subdomain name under Account A before you can connect the `*.example.com` wildcard domain name under Account B.

Does CDN support connecting wildcard domain names?

Yes, CDN supports connecting wildcard domain names, for which domain name ownership verification is required. Once verified, domain names can be connected or retrieved.
In addition:

1. If a wildcard domain name such as `*.test.com` is already connected to Tencent Cloud, then none of its sub-domain names can be connected to another account.
2. If the `*.test.com` wildcard domain name is connected under the current account, wildcard domain names in a format such as `*.path.test.com` can be connected only under the current account.
3. If multiple levels of nested domain names are connected at the same time under an account, the domain name settings are applied and traffic statistics are calculated based on the priorities of the domain names. A more accurate match between the accessed domain name and a connected domain name indicates a higher priority for the connected domain name. For example, the access traffic to `a.path.test.com` is accelerated based on the settings of `a.path.test.com` , and the access traffic to `b.path.test.com` is accelerated based on the settings of `*.path.test.com` .

Why do I get an error that the VOD domain name cannot be accessed?

It's because your domain name has already been added to the VOD console. If you want to manage the domain name in the CDN console, it must be deleted from the VOD console and wait about 1 minute before adding it to the CDN console, or access other subdomain names.

How long does it take to configure CDN?

Most CDN configurations take effect within 5 minutes. Some CDN configurations take effect within 5 to 15 minutes because a large number of tasks need to be run to complete the configurations. Please wait.

Can I configure multiple origin server IPs?

Yes. After you configure multiple IPs, CDN will randomly access one of the IPs when forwarding a request to the origin server. If the number of origin-pull failures with this IP exceeds the threshold, the IP will be isolated for 300 seconds by default, during which no origin-pull requests will be forwarded to the IP.

How do I bind CNAME to a domain name after the domain name is connected to CDN?

See CNAME Configuration for how to bind CNAME with your DNS service provider.

What business types does CDN support?

The selected service type determines which resource platform is used by the domain name. Acceleration configurations vary by resource platforms. Please choose the service type that matches your business:

- Acceleration of small webpage file downloads: applicable to e-commerce, websites, UGC communities, and other business scenarios that mainly involve small static resources, such as webpage styles, images, and small files.
- Acceleration of large file downloads: applicable to business scenarios where large files, such as game installation packages, application updates, and application program packages, are downloaded.
- Audio and video on demand acceleration: applicable to audio and video on-demand scenarios that require acceleration, such as online on-demand audio and video streaming.
- Dynamic and static content acceleration: applicable to business scenarios where dynamic and static data is integrated, such as various website homepages.
- Dynamic content acceleration: applicable to scenarios such as account login, order transactions, API calls, and real-time queries.

Why do exceptions such as old resources, old content, and incorrect content occur after the acceleration?

CDN nodes cache resources based on the cache validity configuration. If the resources that are cached on a CDN node are not expired, the CDN node does not synchronize the latest resources from the origin server.
After a resource on the origin server is updated, its cache on the CDN node must be updated immediately. You can use the cache purge feature to update unexpired caches on the CDN node, so as to ensure that resources cached on the CDN node and stored on the origin server are consistent.

How do I modify the project of a domain name in CDN?

Log in to the CDN console, select **Domain Management** on the left sidebar, click **Manage** on the right of a domain name, open the **Basic Configuration** tab, and then modify the domain name project. To modify the projects of multiple domain names in batches, please tick target domain names on the **Domain Management** page, click **More Actions** drop-down list, and select **Edit Project**. Up to 50 domain names can be operated at a time.

> Note：
> Users on the CDN permission system should proceed with caution, since this operation may cause changes to the permissions of sub-users.

My domain name has already obtained an ICP filing from the MIIT. Why does the system prompt that it does not have an ICP filing when I try to connect it to CDN?

After you obtain your ICP filing, it takes some time to sync the information from the MIIT to Tencent Cloud CDN. Please wait 24 hours and try again.

Can I configure ports for acceleration domain names or origin servers?

- CDN acceleration domain name port: currently, CDN acceleration ports can only be 80, 443, and 8080.

- Origin server port: the ports 1 to 65535 can be configured after the origin server address.

What is CDN origin domain configuration?

The origin domain is the website domain name that is accessed on the origin server during origin-pull on a CDN node. The IP or domain name that is configured on the origin server allows a CDN node to find the corresponding origin server during origin-pull. If multiple websites run on the origin server, the origin domain configuration specifies the domain name of the website to which requests are forwarded. If only one website runs on the origin server, you do not need to modify the origin domain, and the acceleration domain name is used as the origin domain by default.
If you use a Tencent Cloud Object Storage (COS) bucket or a bucket of a third-party object storage service as the origin server, you cannot modify the origin domain, and the origin-pull address is used as the origin domain by default.

How do I tell whether CDN has taken effect?

1. View the domain name list in the CDN console. The CDN acceleration service has taken effect for your domain name if at least one CNAME resolution record is valid for the domain name. This means that the CNAME resolution of the domain name is complete.

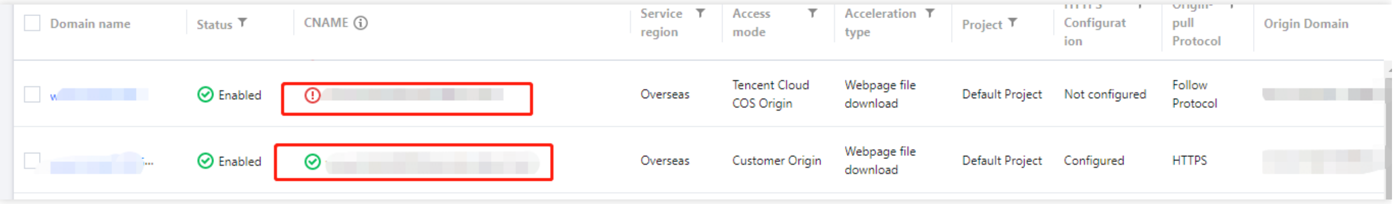

□ Domain name	Status	CNAME ⓘ	Service region	Access mode	Acceleration type	Project	HTTPS Configuration	Origin-pull Protocol	Origin Domain
□ [redacted]	✓ Enabled	ⓘ [redacted]	Overseas	Tencent Cloud COS Origin	Webpage file download	Default Project	Not configured	Follow Protocol	[redacted]
□ [redacted]	✓ Enabled	✓ [redacted]	Overseas	Customer Origin	Webpage file download	Default Project	Configured	HTTPS	[redacted]

2. Alternatively, run the `nslookup` or `dig` command. In this example, the domain name is `www.test.com` .
   - If you use Windows, open the command prompt and run the `nslookup -qt=cname www.test.com` command. Check the CNAME resolution record in the output. If the CNAME resolution record is the same as the CNAME address that is provided by CDN, the CDN acceleration service has taken effect for the domain name.

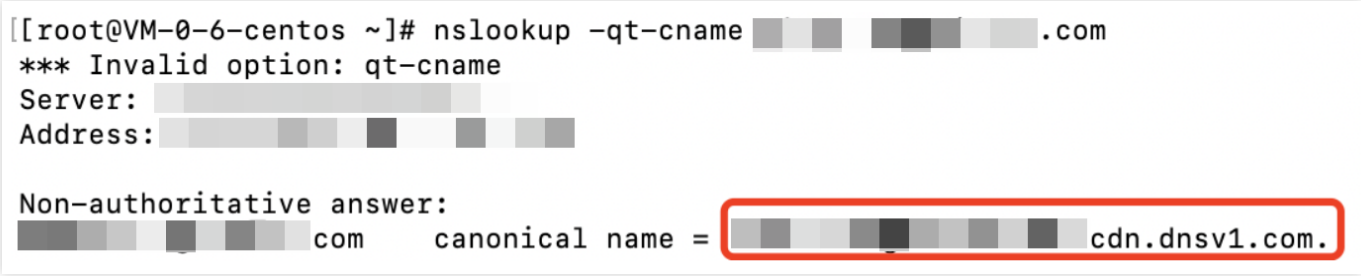

```
[root@VM-0-6-centos ~]# nslookup -qt-cname [redacted].com
*** Invalid option: qt-cname
Server:    [redacted]
Address:   [redacted]

Non-authoritative answer:
[redacted]com    canonical name = [redacted]cdn.dnsv1.com.
```

   - If you use macOS or Linux, open the command prompt and run the `dig www.test.com` command. Check the CNAME resolution record in the output. If the CNAME resolution record is the same as the CNAME address that is provided by CDN, the CDN acceleration service has taken effect for the domain name.

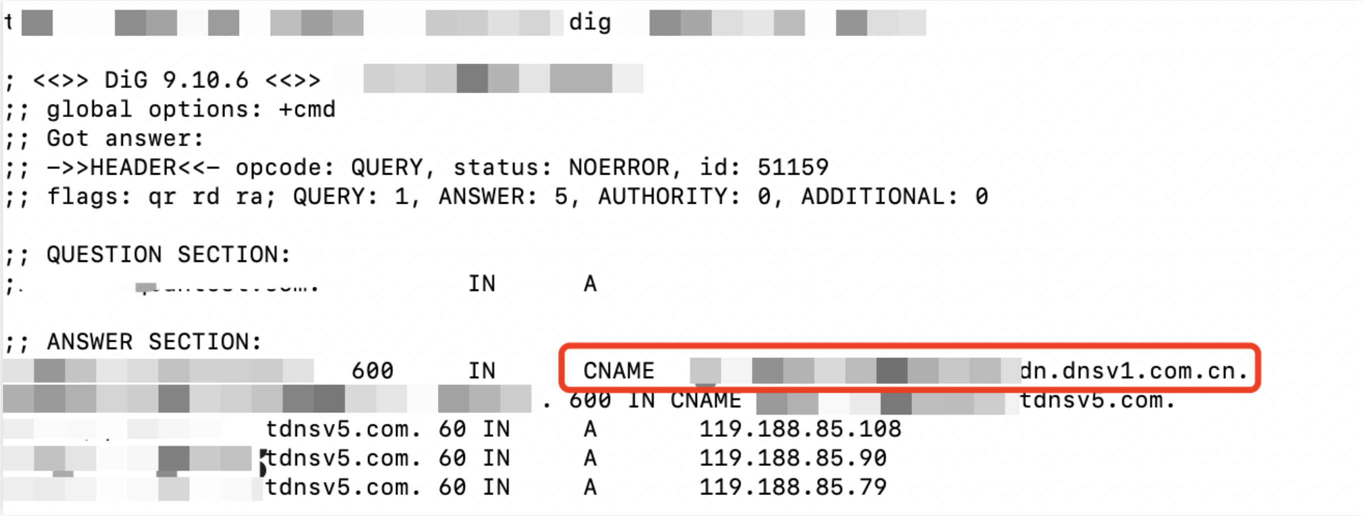

```
t [redacted] dig [redacted]

; <<>> DiG 9.10.6 <<>> [redacted]
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 51159
;; flags: qr rd ra; QUERY: 1, ANSWER: 5, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;[redacted]                    IN      A

;; ANSWER SECTION:
[redacted]          600     IN      CNAME   [redacted]dn.dnsv1.com.cn.
[redacted]. 600 IN CNAME [redacted]tdnsv5.com.
[redacted]tdnsv5.com. 60 IN     A       119.188.85.108
[redacted]tdnsv5.com. 60 IN     A       119.188.85.90
[redacted]tdnsv5.com. 60 IN     A       119.188.85.79
```

What do I do if files fail to be downloaded from CDN?

If files cannot be downloaded from CDN, we recommend troubleshooting by the following methods:

1. Check whether files can be downloaded normally from the origin server.
2. Check whether the domain name is correctly configured in the CDN console (see the origin domain in the **Basic Configuration** tab). Make sure that the configured origin domain name can be connected properly. Otherwise, origin-pull may fail, which will affect your business.
3. Check the security policy of the origin server. Check whether the origin-pull failure is caused by the security policy that is configured on the origin server. If so, obtain the intermediate IP range and add the origin server to the allowlist.

What should I do if I cannot log in to the WordPress backend after CDN acceleration is configured?

When you configure CDN acceleration for WordPress, you must properly configure cache rules for resources that are related to dynamic requests, such as interface-related resources and login-related resources (resources in the /wp-admin backend login address). Otherwise, you may encounter a login failure. We recommend that you do not cache dynamic files.

What do I do if the origin-pull protocol or port is invalid when I configure the origin server?

Tencent Cloud CDN supports port customization when you configure the origin server. If you set the origin-pull protocol to HTTP, port 80 is used for origin-pull by default. If you set the origin-pull protocol to HTTPS, port 443 is used for origin-pull by default. If you configure a custom port, the custom port is used for origin-pull. Make sure that you properly configure the origin-pull protocol and port when you configure the origin server. Otherwise, origin-pull may fail. The following examples list the common configuration errors:

1. The origin-pull protocol is set to HTTP, but the origin server supports only HTTPS-based origin-pulls.
2. The origin-pull protocol is set to HTTP, and the custom port 443 is used. However, the origin server supports only HTTPS-based origin-pulls.
3. The origin-pull protocol is set to HTTP, and the custom port 8080 is used. However, the origin server does not support access requests from port 8080.

If the origin-pull protocol is valid and the default port is invalid, use a custom port. After you enter the information about the origin server, the system automatically checks whether the origin server supports access from the custom port and returns the check result. If the check fails, troubleshoot issues based on the returned check result.

Does CDN support .top domain names?

Yes. CDN already supports domain names suffixed with .pw or .top.

Does CDN support Chinese domain names?

CDN supports domain names that contain underscores (_) and Punycode-converted Chinese characters.

- You must obtain ICP filings for Chinese domain names before you convert the Chinese characters in the domain names to Punycode.
- After you add a Chinese domain name, to the allowlist, you can convert the domain name to `xn--fiq228c.xn--eqrt2g` by using a thrid-party tool, and then connect `xn--fiq228c.xn--eqrt2g` to CDN.
- You can directly add domain names that contain underscores (_), such as `test_qq.tencent.cloud` .

What will happen to the files on CDN nodes if I disable the connected domain name in the CDN console?

If you disable the CDN service of a connected domain name, CDN nodes will retain the connection configurations of the domain name, CDN traffic will no longer be generated, and the domain name will be inaccessible.

Why does the “The CAM policy is not configured for the sub-account” error message appear?

The error message appears when you use a sub-account to perform operations, such as adding domain names and querying data, and you have not used the root account to attach policies to the sub-account. To resolve the problem, use the root account to go to the Policies page in the CAM console, create CDN-related policies, and attach the policies to the sub-account. After the authorization is complete, you can go to the user list to view the policies that are attached to the sub-account.

How do I disable or delete an acceleration domain name? Will the configuration be retained after the acceleration domain name is disabled or deleted?

To stop acceleration, log in to the CDN console, disable the domain name, and then delete the domain name. For more information, see Domain Name Operations. If you cannot delete the domain name after you disable it, check whether the domain name is in the Disabling state. If not, check whether you are logged in to the CDN console with a sub-account. If yes, use the root account to grant the required permissions to the sub-account.
After a domain name is disabled, resources that have been configured are retained, but the acceleration stops and the 404 error code is returned for incoming user requests. After the domain name is deleted, resources that have been configured are immediately deleted and cannot be recovered.

How do I enable CDN acceleration for the `example.com` , `www.example.com` , and `m.example.com` domain names at the same time?

1. To enable CDN acceleration for the three different domain names, connect them one by one to CDN. If you want to apply the same settings to the three domain names, add the domain name in batches or replicate the settings when you add the domain names.
2. To access the same resource from multiple domain names, such as `example.com` and `www.example.com` , add an implicit or explicit URL at your DNS provider to point the domain names to a website by using the 301 redirect technology. For more information, see Implicit and Explicit URL Records.

Does CDN support a WebSocket connection?

We recommend that you enable dynamic and static content acceleration or dynamic content acceleration by using Enterprise Content Delivery Network (ECDN). You can configure the timeout period for the WebSocket connection. The timeout period can be up to 300 seconds. The WebSocket connection may be unstable or even fail when you use the following acceleration types: small webpage file downloads, large file downloads, and audio and video on demand.

Cache Configuration FAQs

Last updated：2022-06-17 16:50:36

What's node cache validity configuration?

Node cache validity configuration refers to a set of validity rules the CDN cache nodes should follow when caching your business contents.
All resources cached on CDN nodes have validity. For unexpired resources, when a request reaches the node, the node will directly return the requested resources to the user, so as to speed up the resource acquisition. For expired resources, the node will forward the user request to the origin server. If the resources have been updated on the origin server, they will be reacquired, cached to the node, and then returned to the user; otherwise, only the resource validity will be updated on the node. A proper cache validity can effectively improve the resource hit rate and lower the origin-pull rate, reducing bandwidth usage.

How do I control the file cache validity in a browser?

You can configure the browser cache validity on the console. For more information, please see Browser Cache Validity Configuration.

How do I configure CDN to return specific files without caching?

You can configure the cache validity for resources based on the directory, file path, file type. For more information, see Node Cache Validity Configuration (New).
When you set "No Cache" for a file, the file will not be cached on CDN nodes. Each time you request the file, CDN nodes will pull it from the origin server directly.

What cache validity configurations supported in CDN?

CDN allows you to set a cache validity period and whether to ignore query string, ignore case, follow origin server and enable heuristic cache for various file types. By using these cache rules properly, you can effectively improve the hit rate with a lower origin-pull rate and bandwidth usage. For details, see Cache Configuration and Node Cache Validity Configuration (New).

What is the default cache configuration of CDN?

When adding an acceleration domain name, default node cache validity rules are added based on different acceleration service types and can be modified as needed.

- The following types of resources are not cached by default, including CDN webpage files, large files and audio and video on demand, and ECDN dynamic and static content (such as PHP, JSP, ASP and ASPX dynamic files). Other files are cached for 30 days.

- For ECDN dynamic content acceleration, all files are not cached.

If no rule is configured or matches requests, the default policies will be applied:

- When a user makes a request for a certain business resource, if the HTTP response header of the origin server contains the field `Cache-Control` , the `Cache-Control` will be followed.
- If the HTTP response header of the origin server does not contain the field `Cache-Control` , then the resource cache validity on nodes will be 600 seconds.

What are cache matching rules?

When multiple cache rules are set, the ones at the bottom of the list have higher priority. For example, if a domain name is configured as follows:

```
All files - 30 days
.php .jsp .aspx - 0 seconds
.jpg .png .gif - 300 seconds
/test/*.jpg - 400 seconds
/test/abc.jpg - 200 seconds
```

If the domain name is `www.test.com` , and the resource is `www.test.com/test/abc.jpg` , the matching rule will be as follows:

1. Match with the first rule. It is hit, so the cache validity is 30 days.
2. Match with the second rule. It is not hit.
3. Match with the third rule. It is hit, so the cache validity is 300 seconds.
4. Match with the fourth rule. It is hit, so the cache validity is 400 seconds.
5. Match with the fifth rule. It is hit, so the cache validity is 200 seconds.

The final cache validity is subject to the last matching result, so it will be 200 seconds.

How do I tell whether user access has hit the CDN cache?

You can check the X-Cache-Lookup of the HTTP response header.

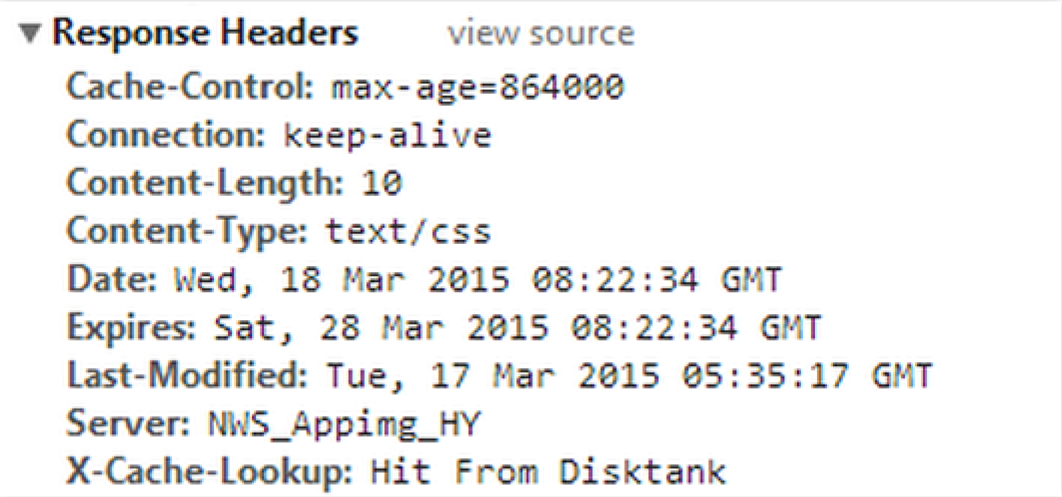

```
▼ Response Headers     view source
    Cache-Control: max-age=864000
    Connection: keep-alive
    Content-Length: 10
    Content-Type: text/css
    Date: Wed, 18 Mar 2015 08:22:34 GMT
    Expires: Sat, 28 Mar 2015 08:22:34 GMT
    Last-Modified: Tue, 17 Mar 2015 05:35:17 GMT
    Server: NWS_Appimg_HY
    X-Cache-Lookup: Hit From Disktank
```

X-Cache-Lookup: Hit From MemCache

X-Cache-Lookup: Hit From Disktank

X-Cache-Lookup: Cache Hit

If any of the above is returned, a cache hit occurs, otherwise it is a cache miss.

If the file changes on the origin server, will the cache on CDN cache nodes be updated in real time?

No. The cache on CDN cache nodes will not be updated in real time.

- CDN cache nodes update the cache according to the cache validity configuration rules you configure in the console. If there are file changes on the origin server and the cache is still valid, CDN cache nodes will not perform origin-pull to update the cache. As a result, the file on the origin server is different from the cache.
- If you need to update the cache of a file, you can purge the cache. When you request the file, CDN will perform origin-pull to get the latest one and re-cache it. You can also prefetch the cache so that CDN pulls the latest resource from the origin server.
- If you need to update the cache of a file regularly, you can enable scheduled purge and prefetch and create a scheduled purge task.

Purge and Prefetch

Last updated：2022-07-22 17:00:23

When do I need to purge and prefetch?

- Purge: To ensure that users access to the latest resources when there are resources to update, restricted resources to remove, or domain name configurations to change on your origin server, you can submit a purge task, which can prevent user access to old resources or old configurations from the node cache. For more details, see Purge Cache.
- Prefetch: For operating activities, installation packages or upgrade packages to release, you can submit a prefetch task to prefetch static resources to CDN acceleration nodes, which will reduce strain on the origin server and improve the service availability and user experience. For more details, see Prefetch Cache.

What are the differences between purge and prefetch?

- Once a resource is purged, its cache on all CDN nodes across the entire network will be deleted. When a user request arrives at a node, the node will pull the corresponding resource from the origin server, return it to the user, and cache it to the node to ensure that the user can obtain the latest resource.
- When a resource is prefetched, it will be cached in advance to all CDN nodes across the entire network. When a user request arrives at a node, the resource can be directly obtained on the node.

What are limits for purge and prefetch? How long do they take to take effect?

- Purge Cache
- URL purge: a maximum of 10,000 URLs can be purged each day and a maximum of 1,000 URLs can be submitted for each purge. It takes about 5 minutes for the purge to take effect. If the cache validity period configured for the file is less than 5 minutes, we recommend you wait for the timeout and update instead of using the purge tool.
- Directory purge: a maximum of 100 directories can be purged each day and a maximum of 500 URL directories can be submitted for each purge. It takes about 5 minutes for the purge to take effect. If the cache validity period configured for the folder is less than 5 minutes, we recommend you wait for the timeout and update instead of using the purge tool.
- Prefetch Resource
- URL prefetch: A maximum of 1,000 URLs can be prefetched each day, and a maximum of 500 URLs can be submitted for each prefetch task. It takes about 5 to 30 minutes for the prefetch to take effect, depending on the file size.

If the resource changes on the origin server, will the cache on CDN cache nodes be updated in real time?

No. The cache on CDN cache nodes will not be updated in real time.

- If the resource changes on the origin server and the cache is still valid, CDN cache nodes will not perform origin-pull to update the cache, and thus the resource on the origin server is different from the cache. In this case, you can specify a proper cache validity period in the console.
- If the cache validity period is too short, CDN will frequently pull the content from the origin server and more traffic will be consumed on the origin server. If it is too long, the cache will be updated slowly.
- If you need to update the cache of a resource, you can purge the cache, and then perform prefetch so that CDN pulls the latest resource from the origin server. You can also send a request to CDN for the latest resource.

How do I view the purge and prefetch history?

You can check the purge and prefetch history in the CDN console. For more information, see History.

Can I prefetch with custom request headers?

No.

How do I increase the daily quota limit for purge and prefetch?

After checking your quota usage at Quota Management in the CDN console, you can request a temporary or permanent quota as needed. The following quota types are supported: URL purge quota, directory purge quota, and URL prefetch quota.

- Temporary quota is a quota that can be applied on a temporary basis and used within a validity period. When it expires, the quota type will end up as permanent.
- Permanent quota is a quota that can be used for an indefinite period. As the permanent quota application takes a long time to process, we recommend requesting a temporary quota to meet your needs.

What should I pay attention to when prefetching?

When you prefetch the file whose cache is still valid, CDN will not perform origin-pull to update the file. We recommend you purge the cache before submitting a prefetch task.

- During prefetching, CDN will pull the required content from the origin server. If you submit a large batch of prefetch tasks, the bandwidth usage of the origin server will greatly increase. Therefore, a proper number of prefetch tasks is suggested.
- CDN provides a dual acceleration structure of edge and middle-layer nodes. If resources are prefetched in the Chinese mainland, they are prefetched to middle-layer nodes in the Chinese mainland, while for those prefetched outside the Chinese mainland, they are prefetched to edge nodes outside the Chinese mainland. Note that traffic incurred when resources are prefetched to edge nodes will be billed.

Statistical Analysis

Last updated：2021-04-06 10:27:58

How are the bandwidth statistics in access monitoring collected?

Each CDN node collects traffic data in real time and reports it to the computing center, which aggregates the data into total traffic data and displays the bandwidth statistics by dividing the total traffic by the duration of use.

Example:

- If the total traffic generated in a minute is 6 MB, then the corresponding bandwidth is (6 * 8) / 60 = 0.8 Mbps.
- As the usage for bill-by-bandwidth is calculated based on the statistics at the 5-minute granularity, the corresponding bandwidth value is total traffic in 5 minutes / 300 seconds.

Why is the traffic in the monitoring information different from that in the log?

The traffic counted based on the downstream bytes in the log of an acceleration domain name is limited to the data at the application layer, while the traffic generated by actual data transfers over the network is around 5-15% more than application-layer traffic.

- Consumption by TCP/IP headers: in TCP/IP-based HTTP requests, each packet has a maximum size of 1,500 bytes and includes TCP and IP headers of 40 bytes, which generate traffic during transfer but cannot be counted by the application layer. The overhead of this part is around 3%.
- TCP retransmission: during normal data transfer over the network, around 3-10% of packets are lost on the internet and retransmitted by the server. This type of traffic cannot be counted by the application layer, which accounts for 3-7% of the total traffic.

As an industry standard, the billable traffic is the sum of the application-layer traffic and the overheads described above. Tencent Cloud CDN takes 10% as the overheads proportion, so the monitored traffic is around 110% of the logged traffic.

How do I calculate the traffic hit rate?

By default, CDN enables L2 cache (edge layer and intermediate layer). As long as a request hits either layer for response, it will be counted as a CDN node hit.

Traffic hit rate = (total downstream traffic - origin-pull traffic) / total downstream traffic.

How do I fix the problem of low traffic hit rate?

- Check whether the cache is purged: cache purge clears the specified content on the node, leading to a temporarily low traffic hit rate.
- Check whether new resources have been put onto the origin server: high numbers of new resources may cause origin-pulls by CDN nodes, leading to a low traffic hit rate.

- Check whether the origin server works properly: if it is malfunctioning with multiple 4XX or 5XX errors, the traffic hit rate will be affected.
- Check whether the cache validity is correctly configured: check the cache validity configuration on the **Cache Configuration** tab in the console. The rules are displayed in ascending order by priority, i.e., a rule takes precedence over the one above it.
- Check whether Range GETs is enabled: check the Range GETs configuration on the **Origin-pull Configuration** tab in the console. If it is disabled, files will be pulled in their entirety instead of multiple parts as requested during origin-pull, which increases the origin-pull traffic and lowers the hit rate.
- Check whether Ignore Query String is enabled: check the Ignore Query String on the **Access Control** tab in the console. If it is disabled, caching will be performed based on the full path. In this case, if the same resource is requested by different parameters, it cannot be matched and will be cached multiple times, lowering the traffic hit rate.

Do status code statistics include all status codes?

Yes. In the new version of CDN statistical analysis, monitoring curves are drawn for all status codes generated by origin servers, making it easier for you to troubleshoot.

How are district and ISP statistics calculated?

The district and ISP statistics are calculated based on the client IPs in the access log. As the calculation is completed based on the log, the simply accumulated billable data differs from the billable data when "all districts" or "all ISPs" is selected. For more information, please see the question 2 above.

How is CDN origin-pull traffic generated?

CDN origin-pull traffic is generated during the following three situations:

1. The requested resource is not cached on the CDN node and is pulled from the origin server.
2. The manually purged origin server is synced with the node.
3. The origin server is automatically purged.

What should I do if my CDN traffic has an exception or is under DDoS or CC attacks?

If you think the number of access requests to your business is moderate, you can download the access logs, and based on which to configure the access limits. CDN does not know your business logic, so no access limits are set for your business by default. Please configure as required by your business.
To avoid malicious requests or CC/DDoS attacks to your website, we strongly recommended you complete the following configurations:

1. Hotlink protection configuration: you can control the access to your business resources. By setting an access control policy on the value of the referer field in the HTTP request header, you can restrict the access source to prevent hotlinking by malicious users. For more information, please see Hotlink Protection Configuration.

2. IP blocklist/allowlist configuration: you can create filtering policies for source IPs of user requests based on your business needs, helping prevent hotlinking and attacks from malicious IPs. For more information, please see IP Blocklist/Allowlist Configuration.
3. IP access limit configuration: you can defend against CC attacks by limiting the number of access requests per second to a node allowed for a client IP. After the configuration is enabled, a 514 error will be returned for requests that exceed the QPS limit. Setting a lower frequency limit may affect the usage of your business by normal high-frequency users. Therefore, please set the threshold according to your actual business conditions and usage. For more information, please see IP Access Limit Configuration.
4. Bandwidth cap configuration: you can configure a bandwidth cap for a domain name. When the bandwidth consumed by the domain name exceeds this cap within a statistical cycle (5 minutes), all access requests will be forwarded to the origin server or the CDN service will be disabled depending on your configuration (in both cases, a 404 error will be returned for all access requests). For more information, please see Bandwidth Cap Configuration.

Is there a delay in using APIs to query data? How long is it?

There is a certain delay in using APIs to query data. Queries of real-time data such as access data and billing data have a delay of around 5-10 minutes, while queries of analytical data such as rankings will have delays of approximately half an hour. The data is calibrated on the backend at around 3 am Beijing Time.

FAQs about HTTPS

Last updated：2021-06-15 14:15:11

What is HTTPS?

HTTPS refers to Hypertext Transfer Protocol Secure, a security protocol that encrypts and transfers data based on the HTTP protocol to ensure the security of data transfer. When configuring HTTPS, you need to provide a certificate for the domain name and deploy it on all CDN nodes to implement encrypted data transfer across the entire network.

Does CDN support HTTPS configuration?

Yes. Tencent Cloud CDN fully supports HTTPS configuration. You can either upload your own certificate for deployment or go to the SSL Certificate Service console to apply for a free third-party certificate provided by TrustAsia.

How do I configure an HTTPS certificate?

You can configure an HTTPS certificate in the CDN console. For detailed directions, please see HTTPS Configuration Guide.

Do the HTTPS certificates on CDN nodes need to be synchronized with HTTPS certificate updates on the origin server?

No. Updating the HTTPS certificate of your origin server does not affect the one configured on CDN. You only need to update the HTTPS certificate on CDN when it is or about to be expired.

Can I deny HTTP access and allow HTTPS access only?

Yes. After successfully configuring an HTTPS certificate, you can use the Forced Redirection feature. HTTP requests from users will be forcibly redirected to HTTPS requests.

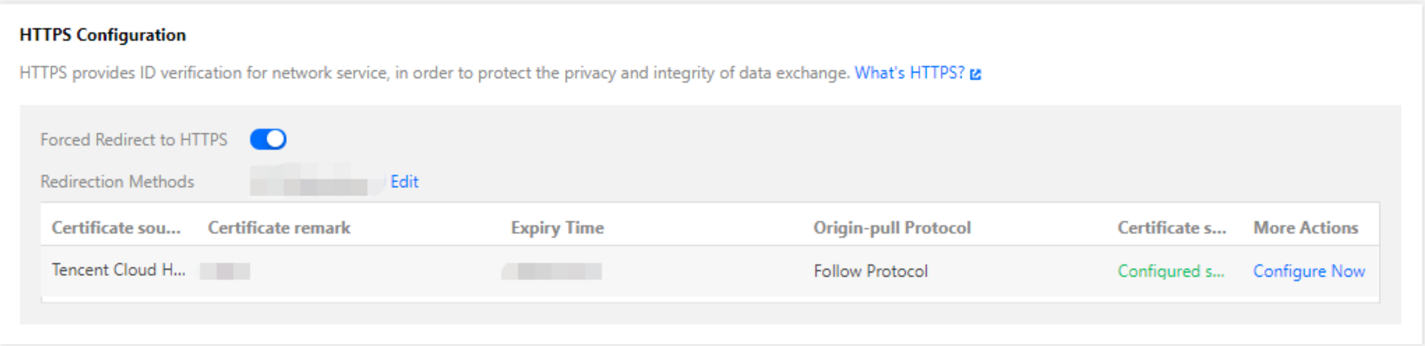

HTTPS Configuration

HTTPS provides ID verification for network service, in order to protect the privacy and integrity of data exchange. What's HTTPS?

Forced Redirect to HTTPS ●

Redirection Methods [redacted] Edit

Certificate sou...	Certificate remark	Expiry Time	Origin-pull Protocol	Certificate s...	More Actions
Tencent Cloud H...	[redacted]	[redacted]	Follow Protocol	Configured s...	Configure Now

Why does HTTPS access not work after CDN is configured?

For HTTPS access, please configure it as instructed:

1. Log in to the CDN console, select **Domain Management** on the left sidebar, and click **Manage** on the right of a domain name to enter its management page.

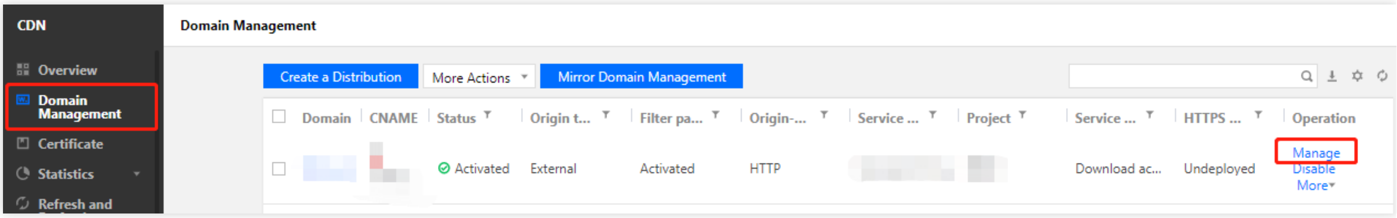

☐	Domain	CNAME	Status	Origin t...	Filter pa...	Origin-...	Service ...	Project	Service ...	HTTPS ...	Operation
☐	[redacted]	[redacted]	Activated	External	Activated	HTTP	[redacted]	[redacted]	Download ac...	Undeployed	Manage Disable More

2. Open the **Advanced Configuration** tab to find the **HTTPS Configuration** section, and click **Configure Now** to go to the **Certificate Management** page. For configuration directions, please see HTTPS Configuration Guide.

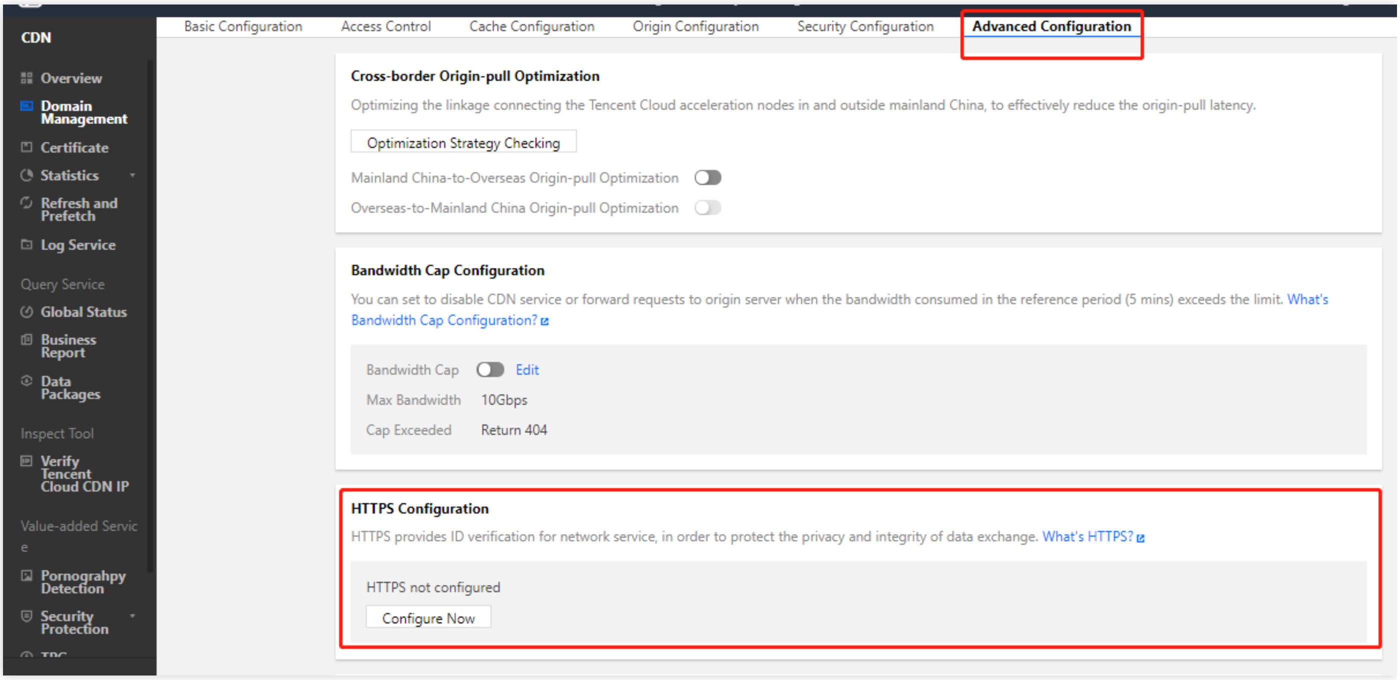

Cross-border Origin-pull Optimization

Optimizing the linkage connecting the Tencent Cloud acceleration nodes in and outside mainland China, to effectively reduce the origin-pull latency.

Optimization Strategy Checking

Mainland China-to-Overseas Origin-pull Optimization

Overseas-to-Mainland China Origin-pull Optimization

Bandwidth Cap Configuration

You can set to disable CDN service or forward requests to origin server when the bandwidth consumed in the reference period (5 mins) exceeds the limit. What's Bandwidth Cap Configuration?

Bandwidth Cap Edit

Max Bandwidth 10Gbps

Cap Exceeded Return 404

HTTPS Configuration

HTTPS provides ID verification for network service, in order to protect the privacy and integrity of data exchange. What's HTTPS?

HTTPS not configured

Configure Now

HTTPS access can be enabled after the certificate is successfully configured.

Connection

Last updated：2023-02-20 15:27:06

How long is the default timeout for Tencent Cloud CDN nodes?

The default timeout is 10 seconds.

What will happen to the files on CDN nodes if I disable the connected domain name in the CDN console?

If you disable the CDN service of a connected domain name, CDN nodes will retain the connection configurations of the domain name, CDN traffic will no longer be generated, and the domain name will be inaccessible.

What should I do if I cannot open my website after connecting it to CDN?

First, as the website cannot be opened when the CDN status of the connected domain name is **Disabled**, please check the status. If the status is not **Disabled**, you can proceed with a check:

- Run `ping` or `nslookup` to check whether the CNAME resolution of the domain name has taken effect. If CNAME has not been added yet, please go to your DNS provider and add CNAME as instructed in CNAME Configuration.
- After CNAME takes effect, you can check whether the origin server is accessible.

If the problem persists, please submit a ticket for assistance.

How do I determine which CDN node is accessed by users?

By running the `nslookup` and `ping` commands, you can get basic troubleshooting information such as the IP, latency, and packet loss of the CDN node accessed by users.

Why do I have a low hit rate?

This is generally due to the following reasons:

- There is a cache configuration problem, such as a short cache validity period.
- HTTP Header prevents caching. Check the origin server's `Cache-Control` or `Expires` configurations.
- There are few cacheable contents for the origin server type.
- The website has few visits and the validity period is short. The low hit rate for files leads to frequent origin-pull requests.

Why do users experience a slow connection when they access CDN?

Please check the download speed for large files and the latency for small files. First, get the URL that is slow to access for users and determine whether the access is slow by using speed test websites such as 17ce (recommended).

If the test result confirms the slow speed and the origin server is an external origin server, please submit a ticket. We will help users check whether the load and bandwidth are restricted on the origin server machine.

How do I tell whether user access has hit the CDN cache?

View the `X-Cache-Lookup` information in the header of the request return. If multiple `X-Cache-Lookup` entries are returned at the same time, that is normal. If `Cache Hit` / `Hit From MemCache` / `Hit From Disktank` is returned, it means that the CDN cache has been hit.

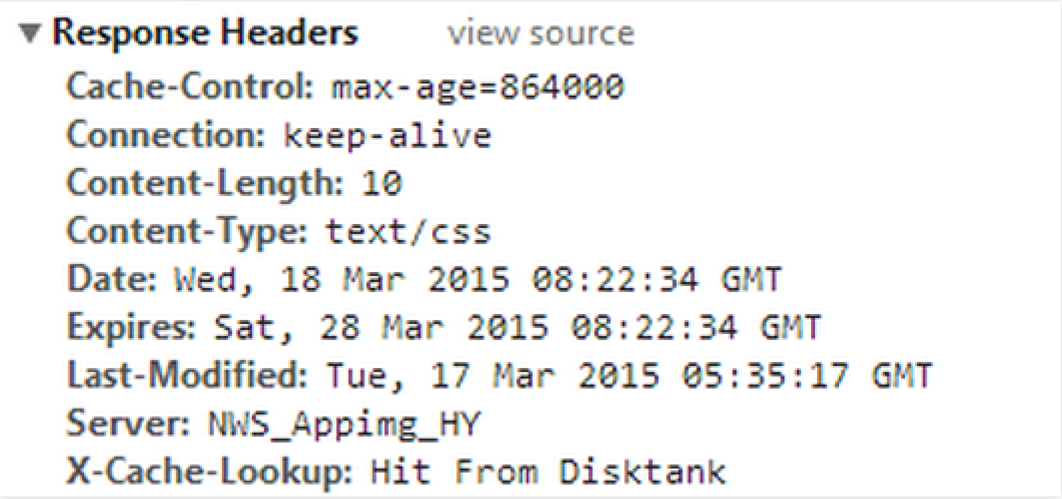

```
▼ Response Headers      view source
    Cache-Control: max-age=864000
    Connection: keep-alive
    Content-Length: 10
    Content-Type: text/css
    Date: Wed, 18 Mar 2015 08:22:34 GMT
    Expires: Sat, 28 Mar 2015 08:22:34 GMT
    Last-Modified: Tue, 17 Mar 2015 05:35:17 GMT
    Server: NWS_Appimg_HY
    X-Cache-Lookup: Hit From Disktank
```

- `X-Cache-Lookup:Hit From MemCache` : the memory of the CDN node has been hit.
- `X-Cache-Lookup:Hit From Disktank` : the disk of the CDN node has been hit.

Why do files with the same name returned by the node have different sizes?

Since all file types are cached by default, there may be different versions of a file on the CDN node. To solve this problem, you can:

- Manually purge files and update the cache immediately.
- Use a version number, e.g., `http://www.xxx.com/xxx.js?version=1` .
- Change the file names to avoid using files with the same name.

If the problem persists, please submit a ticket for assistance.

What should I do if my website cannot be accessed after the hotlink protection allowlist is configured in CDN?

Please select **Allow blank referer** when configuring the hotlink protection allowlist, so that the website can be opened normally in a browser (with a blank referer).

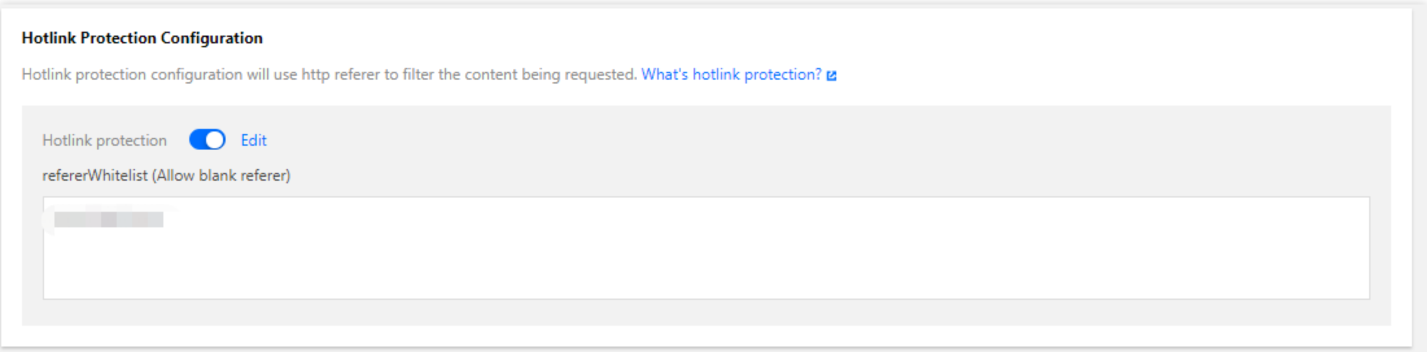

Hotlink Protection Configuration

Hotlink protection configuration will use http referer to filter the content being requested. What's hotlink protection?

Hotlink protection ● Edit

refererWhitelist (Allow blank referer)

[redacted]

Can the traffic cap configuration defend against DDoS attacks?

CDN mainly focuses on content delivery acceleration rather than DDoS protection. You can use CDN’s bandwidth cap feature to automatically collect bandwidth usage statistics in 5 minutes. If the cap threshold is reached, CDN will respond according to the configuration. **The maximum threshold is 10,000 Tbps.** To protect your site against DDoS attacks, use Secure Content Delivery Network.

Can CDN provide all of its node IPs?

No. For security concerns, CDN cannot provide a full node IP list. However, you can query the IP region on the **Verify Tencent Cloud CDN IP** page. For more information, please see Verifying Tencent IPs.

Errors

Last updated：2020-07-23 16:15:12

What should I do if a 423 error is reported in CDN?

The 423 status code is a Tencent Cloud CDN custom status code. CDN will report a 423 error when it detects a loopback request. We recommend you check the following:

1. Check the origin server configured in the CDN Console. If your origin server is also a Tencent Cloud CDN acceleration domain name, it may cause loopback requests.
2. If your origin server is configured with 301/302 HTTP to HTTP redirection and follow 301/302 is enabled in the CDN Console, a 423 error may occur. We recommend you disable follow 301/302.

> ⚠ **Note：**
>
> if you use this method, we recommend you enable HTTPS configuration to force HTTPS redirection and change the origin-pull method to protocol follow. Otherwise, multiple redirects may occur. For the configuration steps, please see HTTPS Acceleration Configuration Guide.

What should I do if a 514 status code is returned in CDN?

This is caused by the IP access limit configured in the CDN Console as shown below:

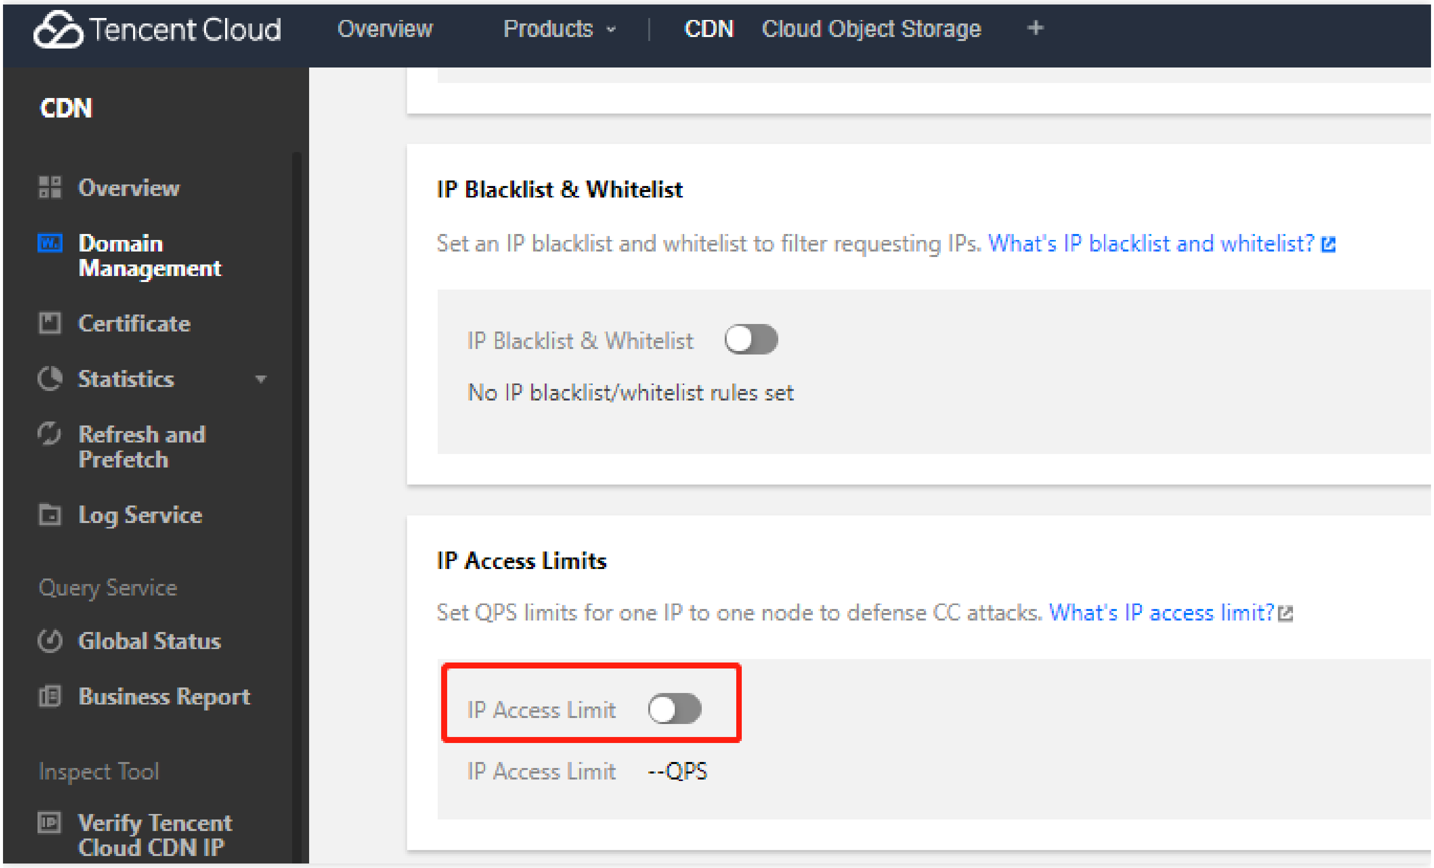

IP Blacklist & Whitelist

Set an IP blacklist and whitelist to filter requesting IPs. What's IP blacklist and whitelist?

IP Blacklist & Whitelist

No IP blacklist/whitelist rules set

IP Access Limits

Set QPS limits for one IP to one node to defense CC attacks. What's IP access limit?

IP Access Limit

IP Access Limit --QPS

Note：

- IP access limit is designed to restrict the number of times an IP is allowed to access a node within one second. If the limit is exceeded, a 514 error will be returned.
- Setting a low limit may affect the use of your business by normal high-frequency users. Therefore, please set a reasonable threshold. For more information, please see IP Access Limit Configuration.

What should I do if a 404 status code is returned by a CDN domain name?

Check the following:

1. Check whether the origin server can be accessed normally.
2. Check whether the origin server information and the origin domain have been modified in the CDN Console. This may have caused the 404 error during origin-pull.

Does CDN support switching to non-origin-pull automatically and return the content cached on a node in case of origin server exception?

In case of origin server exception, if the cache on the CDN node accessed by the user request has not expired, the node will directly return the requested content to the client. If the cache has expired, it will not respond and an origin-pull request will be initiated.