Cloud Load Balancer

FAQs

Product Documentation

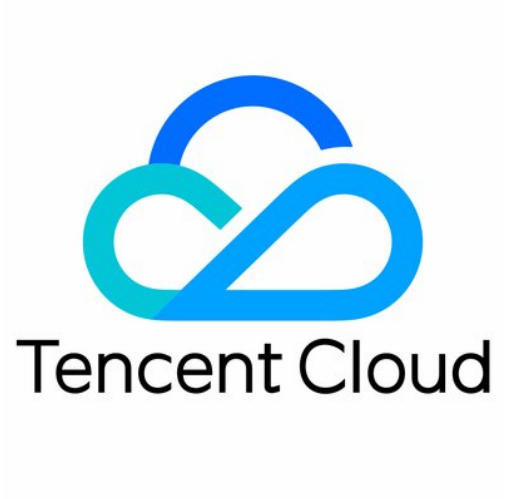

Contents

FAQs
Billing

Last updated：2024-01-04 14:39:00

FAQs About Billing

Do CLB instances and backend CVM instances communicate over the public network or private network?

Communication between CLB and CVM instances generates private network traffic, irrespective of whether the account type is bill-by-IP or bill-by-CVM. The traffic routes of the two types are as follows:
A client accesses a CLB instance over the public network, the CLB instance forwards the traffic to a CVM instance over the private network, then the CVM instance responds to the CLB instance over the private network, finally the CLB instance responds to the client.
When a CVM instance accesses the public network or is accessed via a public IP address, the CVM instance interacts with a client via a public IP address or EIP.

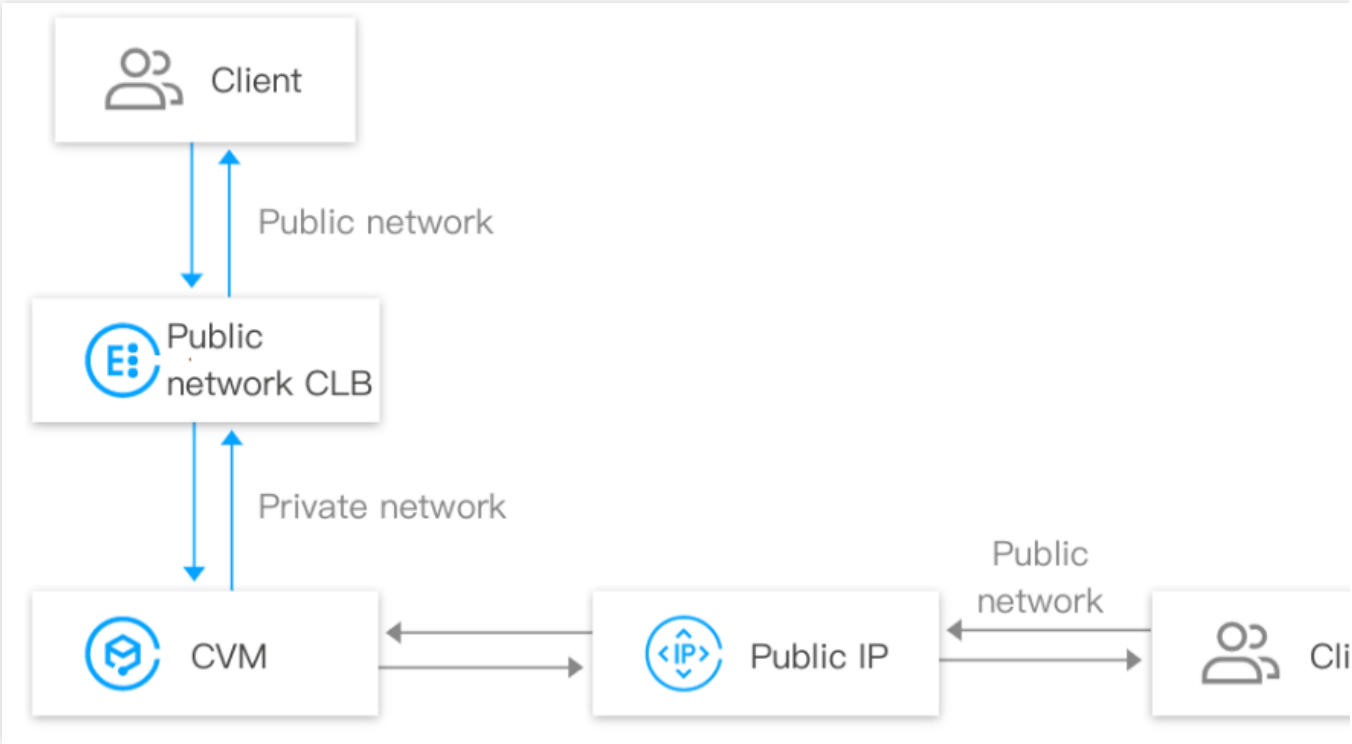

[Back to Top]

How is CLB billed?

CLB billing policies vary by account type. For more information, see Billing Overview. The main difference between bill-by-IP and bill-by-CVM accounts lays in the billing dimension, and other aspects such as service access and monitoring are the same. For more information about differences between account types, see Checking Account Type.
Public network fee of bill-by-IP accounts: For the public network traffic accessing a CLB instance or a public IP address, the public network fee is charged on the CLB instance or the public IP address.
Public network fee of bill-by CVM accounts: For the public network traffic accessing a CLB instance or a public IP address, the public network fee is charged on the CVM instance.

[Back to Top]

Can I switch my account type between bill-by-IP and bill-by-CVM?

You can upgrade your account type from bill-by-CVM to bill-by-IP.
A bill-by-IP account cannot be downgraded to a bill-by-CVM account.

[Back to Top]

How is cross-region binding billed?

The cross-region binding fee is charged only if you use the CLB cross-region binding feature.
Cross-region 1.0: The cross-region binding fee is charged on the CLB instance.
Cross-region 2.0: The cross-region binding fee is charged on the CCN instance, rather than the CLB instance.

[Back to Top]

What is the bandwidth cap of a CLB instance?

Bill-by-IP account: You need to specify the CLB bandwidth cap when you purchase a CLB instance. If the business bandwidth exceeds the cap, the CLB instance will lose packets automatically.

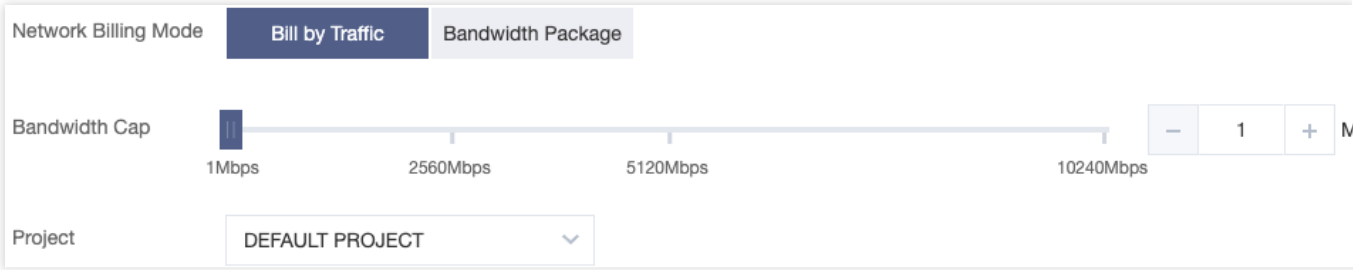

Bill-by-CVM account: You do not need to specify the CLB bandwidth cap when you purchase a CLB instance. The CLB bandwidth is restricted by the total public network bandwidth of the backend CVM instances.

[Back to Top]

What are the monitoring metrics related to billing for CLB?

You can refer to the following monitoring metrics for billing related to CLB, including network fees and LCU fees:

Type	Metric	Description
Network fee	AccOuttraffic	CLB-to-real server traffic out
	OutTraffic	CLB-to-real server bandwidth out
	InTraffic	CLB-to-real server bandwidth in
LCU fee	ClientConnum	Client-to-CLB connections
	ClientNewConn	Client-to-CLB new connections

	TotalReq	Requests per second
	ClientAccOuttraffic	Client-to-CLB traffic out
	ClientAccIntraffic	Client-to-CLB traffic in

You can refer to the following monitoring metrics related to CLB limits, including bandwidth and LCU caps:

Type	Metric	Description
Bandwidth cap	ClientOuttraffic	Client-to-CLB bandwidth out
	ClientIntraffic	Client-to-CLB bandwidth in
LCU cap	ClientConcurConn	Client-to-CLB concurrent connections
	ClientNewConn	Client-to-CLB new connections
	TotalReq	Requests per second
	ClientOuttraffic	Client-to-CLB bandwidth out
	ClientIntraffic	Client-to-CLB bandwidth in

For more information about monitoring metrics, see Monitoring Metrics. For more information about network fees, see Network Fee. For more information about LCU billing, see LCU Pricing.

[Back to Top]

CLB Configuration

Last updated：2024-01-04 14:39:00

Concepts

What is the difference between layer-4 and layer-7 load balancing?

What is the difference between the UDP and TCP protocols?

How does CLB achieve session persistence based on cookies?

What is the real server weight?

What is the difference between resetting the weight to 0 and unbinding the real server?

Health Check

What should I do if the health check result is abnormal?

Why is the frequency of health checks so high?

Access

HTTP redirection in load forwarding

Which TCP ports can CLB be performed on?

What can I do if no policy file is returned and the connection is interrupted after a policy request (i.e., flash server request) is sent from port 843?

Can CLB instances directly obtain client IPs?

Can I configure a private CLB for a CVM instance to forward traffic from port A to another port on the same server?

Does a backend CVM instance need public network bandwidth? Will the bandwidth affect the CLB service?

Notes on compatible versions if the client and server have different HTTP versions

Gzip compatibility

How to configure the security group of a CLB real server? How to configure the access blocklist?

Configuring the access blocklist

Does a CLB instance communicate with its real server over the private or public network?

Pinging CLB VIPs

Running `Telnet` command to connect CLB listening ports

Private network loopback

When a client accesses the same port of a real server via different intermediate nodes, these nodes cannot be specifically determined.

I have configured a backend CVM instance with a security group to block access traffic from the public network and only allow that from CLB instances. But it does not take effect.

I have configured a CLB instance with a listener and bound the listener to a backend CVM, and resolved a domain name to the IP of the backend CVM. But when the domain name is accessed, no monitor data are displayed.

I have not created an 843 listener, but I can successfully run the command `Telnet` for connection.

Is the 9/11 IP range recorded in CLB access logs a private IP range of Tencent Cloud private network?

What is the difference between layer-4 and layer-7 load balancing?

Layer-4 load balancing is based on IPs and ports.
Layer-7 load balancing is based on application layer information such as HTTP headers and URLs.
The difference between layer-4 and layer-7 instances is whether layer-4 or layer-7 information is used as the basis for determining how to forward traffic for load balancing on real servers.
For example, a layer-4 CLB instance determines which traffic needs load balancing based on the layer-3 IP address (VIP) and layer-4 port number. It performs Network Address Translation (NAT) on the traffic to be processed, and then forwards it to the real server. It also records which server has processed the TCP or UDP traffic, and forwards all subsequent traffic of this connection to the same server for processing.
A layer-7 CLB is layer-4 CLB instance combined with application layer features.
For example, CLB instances of the same web server can identify traffic that needs to be processed based on layer-7 URL, browser type, and language in addition to the VIP and port 80.
Layer-7 load balancing is also known as "content exchange", in which an internal server is selected based on meaningful application layer content in the message and the server selection method configured on the CLB instance. To select the server based on the real application layer content, a layer-7 CLB instance must establish a connection (three-way handshake) with the client as a proxy of the final server to receive the message containing real application layer content from the client. It then selects the internal server according to the specific fields in the message and the server selection method configured on the CLB instance. In this case, the CLB instance is more like a proxy server, establishing a TCP connection with the frontend client and the real server respectively.

[Back to Top]

What is the difference between UDP and TCP protocols?

TCP is a connection-oriented protocol. Before receiving and sending data, TCP requires a reliable connection to be established with the other side. UDP is a message-oriented (connection-less) protocol. It directly sends data packets without performing a three-way handshake with the other side. UDP is suitable for scenarios that focus more on real-timeliness than reliability, such as video chat, real-time push of financial market information, DNS, and IoT.

[Back to Top]

How does CLB achieve session persistence based on cookies?

In cookie insertion mode, the CLB instance inserts cookies and the real server does not need to make any modifications. When you make the first HTTP request, it (without a cookie) enters the CLB instance, which then selects a real server based on the load balancing algorithm policy and sends the request to the server. The real server then sends an HTTP response (without a cookie) that is sent back to the CLB instance, which then inserts the cookie and returns the HTTP response (with a cookie) to the client.

When you make the second HTTP request (with the cookie inserted by the CLB instance last time), it enters the CLB instance, which then reads the session persistence values in the cookie and sends the request (with the same cookie as above) to the specified real server. The server gives an HTTP response. Because the server does not write the cookie, the response does not contain the cookie. When the response traffic re-enters the CLB instance, CLB will write the updated session persistence cookie into the response.

[Back to Top]

What is the real server weight?

You can specify the forwarding weight for each CVM instance in the real server pool, and the instance with a higher weight will be assigned with more access requests. You can configure the weights of the backend CVM instances based on their service capabilities and statuses.
If you have also enabled session persistence, the same access request may be forwarded to different real servers. We recommend temporarily disabling session persistence and then checking whether the problem persists.

[Back to Top]

What is the difference between resetting the weight to 0 and unbinding the real server?

Resetting the weight to 0: TCP listeners keep forwarding existing connections, UDP listeners keep forwarding connections with the same quintuple, and HTTP/HTTPS listeners keep forwarding existing connections.New connections on TCP, UDP, HTTP/HTTPS listeners will no longer be forwarded to RS with a weight of zero.
Unbinding the real server: TCP/UDP listeners cease forwarding residual connections instantaneously, while HTTP/HTTPS listeners persist in forwarding residual connections. Upon completion of forwarding residual connections, the connection with RS is severed.

[Back to Top]

What should I do if the health check result is abnormal?

Please troubleshoot by following the steps:
Ensure that you access your application service directly via the real server.
Ensure that the relevant port is open on the real server.
Check whether there is any security software like a firewall on the real server. This may cause the CLB instance to be unable to communicate with the real server.
Check whether the health check parameters of the CLB instance are configured correctly.
We recommend that you use static pages for health check.

Check whether there is high load on the real server that leads to slow response.
Ensure that there are no IP restrictions (`iptables`) on the real server.

[Back to Top]

Why is the frequency of health checks so high?

Suppose that you configured to send a health check packet every 5 seconds in the console. However, the real server receives one or even more health check requests in 1 second.
Highly frequent health check relates to the implementation mechanism of CLB health check. Suppose that 1 million requests from clients are distributed to 4 CLB servers before being sent to real servers, and each CLB server conducts health checks separately. If the CLB servers are configured to send a health check request every 5 seconds, each CLB server will send a health check request every 5 seconds. This is why the real server receives multiple health check requests. For example, if there are 8 CLB servers in a cluster and each sends a request every 5 seconds, then the real server may receive 8 health check requests in 5 seconds.
The advantages of this implementation scheme are high efficiency, accurate check, and avoidance of mistaken removal. For example, if one of the CLB servers in the cluster fails, the other 7 servers can still forward traffic normally. Therefore, if your real server is checked too frequently, you can configure a longer check interval (for example, 15 seconds).

[Back to Top]

HTTP redirection in load forwarding

When you visit the website `http://example.com` via a browser, a redirection to the root directory is required for the server. When you visit the website `http://example.com/` via a browser, the server will directly return the default page of the website's root directory. Similarly, if `http://cloud.tencent.com/movie` is redirected to `http://cloud.tencent.com/movie/` through URL rewriting, entering `http://cloud.tencent.com/movie` will result in an additional URL rewriting process, leading to slight performance degradation and time consumption. However, if `http://cloud.tencent.com/product` is redirected to a page other than `http://cloud.tencent.com/product/` through URL rewriting, you need to consider whether to add "/" after the secondary URL.
In Tencent Cloud CLB, if the frontend and backend port numbers are different, "/" needs to be added after the secondary URL upon the visit to the secondary page, so as to avoid port number changes after HTTP redirection and ensure normal page access.
Assume that in layer-7 forwarding, port 80 on the CLB instance and port 8081 on the real server are listened to. If the client accesses `http://www.example.com/movie` , the access request is forwarded to the real server via the

CLB instance, and the server redirects the request to `http://www.example.com:8081/movie/` (listening port is 8081). In this case, the client access fails (port error).
Therefore, we recommend rewriting the access request to a secondary URL with "/", such as `http://www.example.com/movie/`. This can avoid HTTP redirection, eliminate unnecessary judgment, and reduce unwanted load. If HTTP redirection is required, make sure that the CLB listening port is the same as that of the real server.

[Back to Top]

Which TCP ports can CLB be performed on?

You can perform load balancing for the following TCP ports: 21 (FTP), 25 (SMTP), 80 (HTTP), 443 (HTTPS), 1024–65535, etc.

[Back to Top]

What can I do if no policy file is returned and the connection is interrupted after a policy request (i.e., flash server request) is sent from port 843?

When the CLB instance receives the policy request from port 843, it will return the general cross-domain policy configuration file. If no policy file is returned and the connection is directly closed, the flash server request may be incorrect.
Ensure that the flash server request sent ends with `\\0`.
Note:
It must end with `\\0` and contain 23 bytes. `\\0` indicates a character whose ASCII code is 0 and only occupies 1 byte.
The normal result returned by port 843 is as shown below:

```
VM_02_sles10_64:/ # perl -e 'printf "<policy-file-request/>%c",0' | netcat -i 1 101.226.62.63 8
<?xml version="1.0"?>
<!DOCTYPE cross-domain-policy SYSTEM "/xml/dtds/cross-domain-policy.dtd">

<!-- Policy file for xmlsocket://socks.example.com -->
<cross-domain-policy>

   <!-- This is a master socket policy file -->
   <!-- No other socket policies on the host will be permitted -->
   <site-control permitted-cross-domain-policies="master-only"/>

   <!-- Instead of setting to-ports="*", administrator's can use ranges and commas -->
   <!-- This will allow access to ports 123, 456, 457 and 458 -->
   <allow-access-from domain="*" to-ports="*" />

</cross-domain-policy>
```

[Back to Top]

Can CLB instances directly obtain client IPs?

IPv6 NAT64 CLB instances do not support obtaining client IPs.
Public network IPv6 and IPv4 layer-7 CLB instances use the `X-Forwarded-For` header to get real client IPs. Acquisition of client IPs is enabled on CLB instances by default but needs to be configured on real servers. For more information, see Obtaining Real Client IPs Over IPv4 CLBs.
Public network IPv6 and IPv4 layer-4 CLB instances (over TCP) can directly get real client IPs on backend CVM instances, and no additional configuration is required. For the private network layer-4 CLB instances purchased after October 24, 2016, the Source Network Address Translation (SNAT) is not conducted. They can directly get real client IPs from servers with no additional configuration.

[Back to Top]

Can I configure a private CLB for a CVM instance to forward traffic from port A to another port on the same server?

No. To access port a on server A (10.66.*.101), the request can be forwarded to port b on server B (10.66.*.102) via a private network CLB instance. But it cannot be forwarded to port b on the same server A (10.66.*.101).

[Back to Top]

Does a backend CVM instance need public network bandwidth? Will the bandwidth affect the CLB service?

Public network bandwidth is not required for the backend CVM instances bound to the CLB instances of any bill-by-IP accounts.
No traffic or bandwidth fee is charged for the CLB instances of any bill-by-CVM accounts. Any public network traffic fees generated by CLB service will be charged on the bound backend CVM instances. We recommend choosing bill-by-traffic for public network bandwidth when purchasing the CVM instance and configuring a reasonable peak bandwidth threshold, so that you do not need to keep track of the fluctuation in total CLB egress traffic. The traffic of internet web business fluctuates considerably and cannot be predicted accurately. When bill-by-bandwidth is selected, it is not cost-effective to purchase excessive bandwidth, yet packet loss may occur during business peaks if insufficient bandwidth is purchased.

[Back to Top]

Notes on compatible versions if the client and server have different HTTP versions

Forwarding compatibility

On the frontend (client), HTTP/1.0 and HTTP/1.1 and lower versions are supported.
On the backend (server), Tencent Cloud uses the HTTP/1.0; HTTP/1.0 and HTTP/1.1 and lower versions are supported.

Note:

HTTP/2 is only supported in HTTPS but not in HTTP, and backward compatibility is allowed on both the client and server.

Gzip compatibility

On the frontend (client), Gzip is supported by HTTP/1.0, HTTP/1.1 and lower versions. Additional configuration is not needed because mainstream browsers all support Gzip.
On the backend (server), because HTTP/1.1 is supported on the CVM instance over Tencent Cloud private network, you don not need to make any configuration, and can directly use HTTP/1.1 configured in Nginx by default to achieve compatibility.

Note:

HTTP/2 is only supported in HTTPS, but Gzip can be used in any HTTP version supported by Tencent Cloud.

[Back to Top]

How to configure the security group of a CLB real server? How to configure the access blocklist?

Configuring the CLB security group

If security group rules have been configured for a real server, the CLB instance may not be able to communicate with the server. Therefore, in layer-4 and layer-7 forwarding, we recommend configuring the security group of a real server as allowing all access requests. If the security group is enabled and accesses from any protocols or IP ranges are allowed by default, IPs of all clients need to be configured to the security group rules of the server IP.
For some malicious IPs, you can add them to the top rules of the security group to prevent them from accessing the real server. You can then allow access requests from all IPs (0.0.0.0) to the local service port, so normal clients can access the server. Security group rules are arranged in priority order and matched from top to bottom.
If health check has been configured for layer-7 CLB forwarding in VPC, in the real server security rules, the CLB VIP must be allowed. Otherwise, the health check may fail.

Configuring the access blocklist

If you need to configure a blocklist for some client IPs to deny their access requests, you can configure the security group associated with the cloud services. The security group rules need to be configured as follows:

Note:

Follow the steps below strictly in the given order, otherwise the blocklist configuration may fail.

Add the client IP and port to be rejected into the security group, and select the option in the policy column to reject access from this IP.

Add another security group rule after completing the above configuration to allow access requests to the port from all IPs by default.

When the configuration completes, the security group rules are as follows:

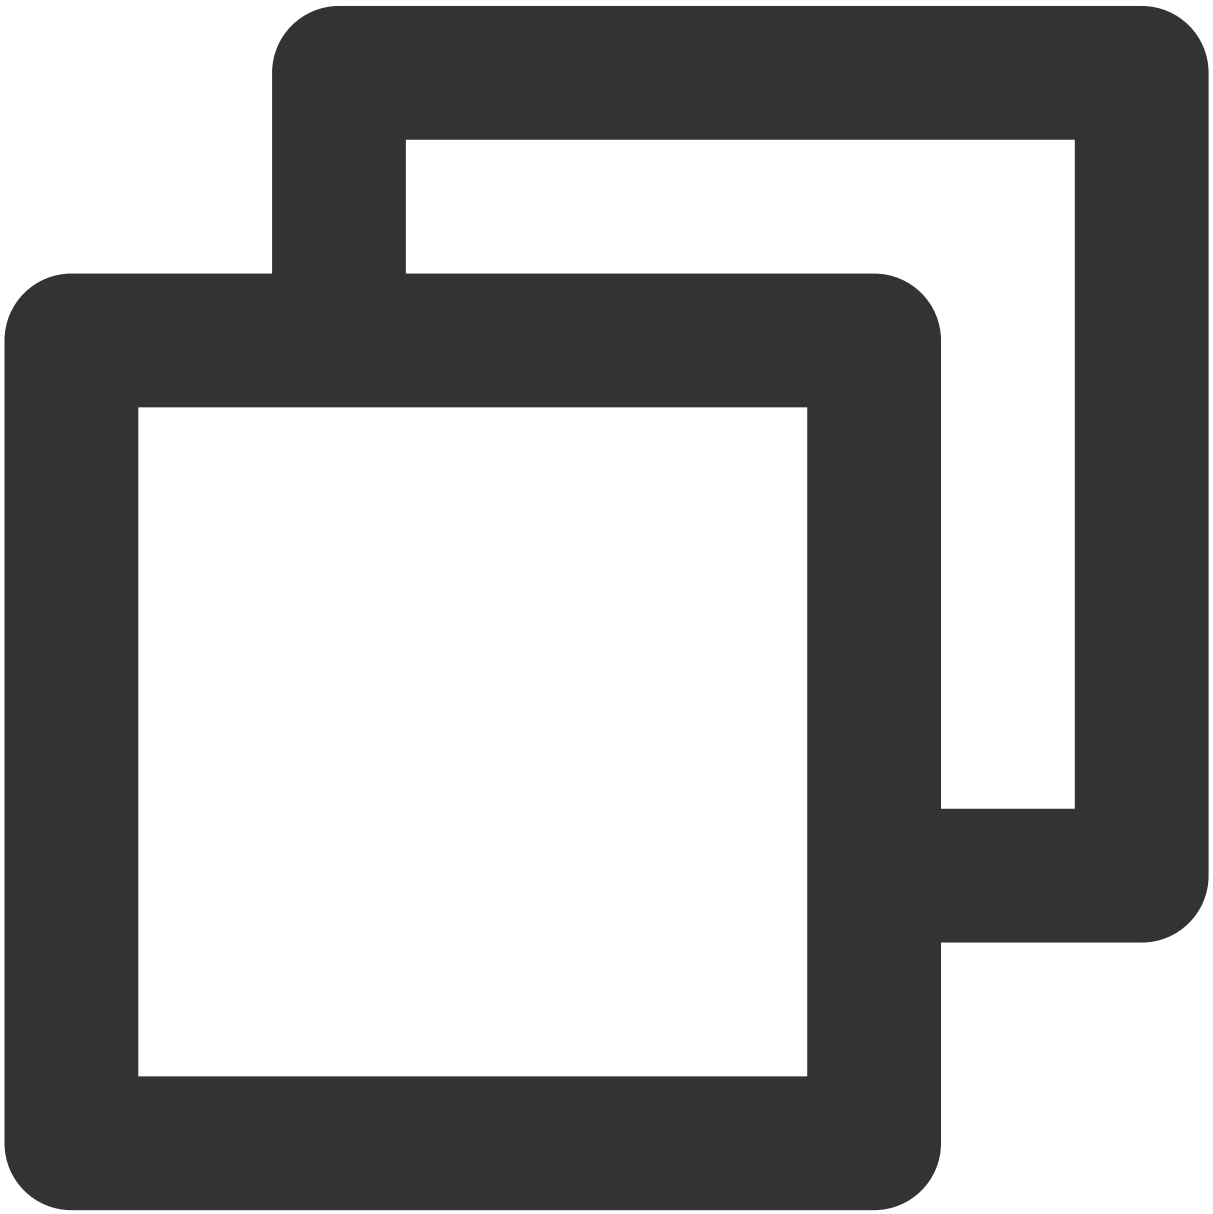

```
clientA ip+port drop
clientB ip+port drop
0.0.0.0/0+port accept
```

For more information about the security group, see Configuring CVM Security Groups.

[Back to Top]

Does a CLB instance communicate with its real server over the private or public network?

A CLB instance always communicates with real servers over the private network, even when the bound CVM instances have public IPs.

[Back to Top]

Pinging CLB VIPs

The requests of pinging the CLB VIP are responded to by the CLB cluster and will not be forwarded to real servers.
A public network CLB VIP can be pinged.
A private network CLB VIP can be pinged only from a client in the same VPC. A ping from a client in another VPC or a local IDC may not succeed because the request is responded by the CLB cluster, which does not reflect the real link. For example, if VPCs are connected via Cloud Connect Network (CCN) or peering connection, we recommend you run the `Telnet` command to test the connectivity to a private network CLB instance.

[Back to Top]

Running `Telnet` command to connect CLB listening ports

If layer-4 listeners (TCP, UDP, and TCP SSL) are not bound with real servers after being created, running the `Telnet` command to connect to their listening ports will fail. It will succeed only if they are bound with real servers. Running the `Telnet` command to connect to listening ports will succeed for layer-7 listeners (HTTP and HTTPS) not bound with real servers, as the CLB cluster will respond instead.

[Back to Top]

Private network loopback

For private network CLB instances, a CVM cannot be both the client and server. When the CLB instances read the same client and server IPs, access will fail.

If your client needs to be used as a server, please bind at least 2 real servers. CLB has related policies to prevent automatic loopback. When client A accesses the CLB instance, the CLB instance will automatically schedule the request to a real server other than client A.

[Back to Top]

When a client accesses the same port of a real server via different intermediate nodes, these nodes cannot be specifically determined.

Issue

When a client accesses the same port of a real server via different intermediate nodes at the same time, these nodes cannot be specifically determined. The scenarios are detailed below:

A client accesses the same port of a real server via layer-4 and layer-7 listeners of the same CLB instance at the same time.

A client accesses the same port of a real server via different listeners of different CLB instances at the same time.

It is more complex to determine intermediate nodes for a client accessing private network CLB instances of a real server rather than multiple clients accessing public network CLB instances of a real server.

Cause

CLB instances pass the client IP to the real server, and `client_ip:client_port -> vip:vport -> rs_ip:rs_port` will change to `client_ip:client_port --> rs_ip:rs_port` .

Solutions

Distributed clients: Multiple clients are used to initiate access.

Fewer CLB instances: The number of CLB instances and listeners are cut down on the premise of meeting business functions and disaster recovery requirements.

Distributed real server ports: Multiple ports for a real server are used to provide services to avoid congestion.

Distributed deployment: Different CLB instances are bound to different ports on real servers. For example, CLB instance 1 bound to one set of CVM instances and CLB instance 2 bound to another set can be accessed at the same time.

[Back to Top]

I have configured a backend CVM instance with a security group to block access traffic from the public network and only allow that from CLB instances. But it does not take effect.

To allow traffic to access a backend CVM instance passing through a CLB instance, the traffic from the public network needs to be allowed both on the backend CVM and CLB security groups. We recommend only allowing the access traffic from a CLB VIP over the public network on the backend CVM security group first, and then allowing public network access IPs on the CLB security group as needed.

[Back to Top]

I have configured a CLB instance with a listener and bound the listener to a backend CVM, and resolved a domain name to the IP of the backend CVM. But when the domain name is accessed, no monitor data are displayed.

Only the traffic going through CLB will be monitored. You can resolve the domain name to the VIP of the CLB instance and access the domain name to view the CLB monitoring information.

[Back to Top]

I have not created an 843 listener, but I can successfully run the command `Telnet` for connection.

The port 843 is opened by default for users to reset the port for Flash access. If you want to close it, you can just do so by not binding any real servers after configuring the `TCP:843` listener.

[Back to Top]

Is the 9/11 IP range recorded in CLB access logs a private IP range of Tencent Cloud private network?

Yes.

[Back to Top]

Troubleshooting Health Check Issues

Last updated：2024-01-04 14:39:00

Tencent Cloud Load Balancer (CLB) instances determine the availability of real servers by performing health checks. This document describes methods for troubleshooting the exceptions detected in a health check.

Note:

If an exception is detected in the health check on a real server, CLB instances stop forwarding traffic to the real server. If all the real servers are checked as abnormal, requests will be forwarded to all the real servers.

For more information about how health check works, see Health Check Overview.

Checking the Public Network Bandwidth of Real Servers

If you use a bill-by-CVM account, you must specify the public network bandwidth for the backend CVM instances bound to CLB instances. Otherwise, the health check result will be abnormal. This is because the bandwidth attribute of the account is on CVM instances instead of CLB instances.

If you use a bill-by-IP account, you do not need to specify the public network bandwidth for the backend CVM instances bound to CLB instances and this does not affect the load balancing service.

Note:

For more information about how to determine your account type, see Checking Account Type.

A bill-by-CVM account does not pay traffic or bandwidth fees. Fees for public network traffic incurred by CLB are billed in the bound backend CVM instances.

You can purchase public network bandwidth for a CVM instance without assigning a public network IP address to it.

Checking Security Group Configuration

Check whether the **Allow Traffic by Default** feature is enabled in the security group of the CLB instance. If not, you need to allow source IP addresses in the CVM security group. If your CLB service supports access from any IP addresses, set the source IP address to `0.0.0.0/0` in the inbound rules of the security group. For more information, see Configuring CLB Security Group.

Checking Layer-4 Listeners

Note:

CLB checks server health by using SYN packets over Transmission Control Protocol (TCP).

CLB checks server health by running the `ping` command over User Datagram Protocol (UDP).

If you find an exception when viewing the health status of the CLB server ports on the page, use the following methods for troubleshooting:
Confirm whether a security group is configured for CLB real servers and thus affecting the service. The security group performs access control on real servers to ensure normal service. For more information, see Configuring CVM Security Groups.
Run the `netstat` command to check whether there is a listening process on the real server port. If no, restart the service.

Checking Layer-7 Listeners

For Layer-7 (HTTP protocol) services, when the health check result on a listener is abnormal, perform the following steps for troubleshooting:
Because the layer-7 health check service of CLB communicates with backend CVM instances over the private network, you need to log in to the server to check whether the application server port is listened on normally at the private network address. If not, enable the listening of the application server port on the private network to ensure the normal communication between the CLB instance and backend CVM instances.
Suppose that both CLB's frontend port and CVM's backend port are 80, and the CVM's private IP is `1.1.1.10` :
For a Windows server, use the following command:

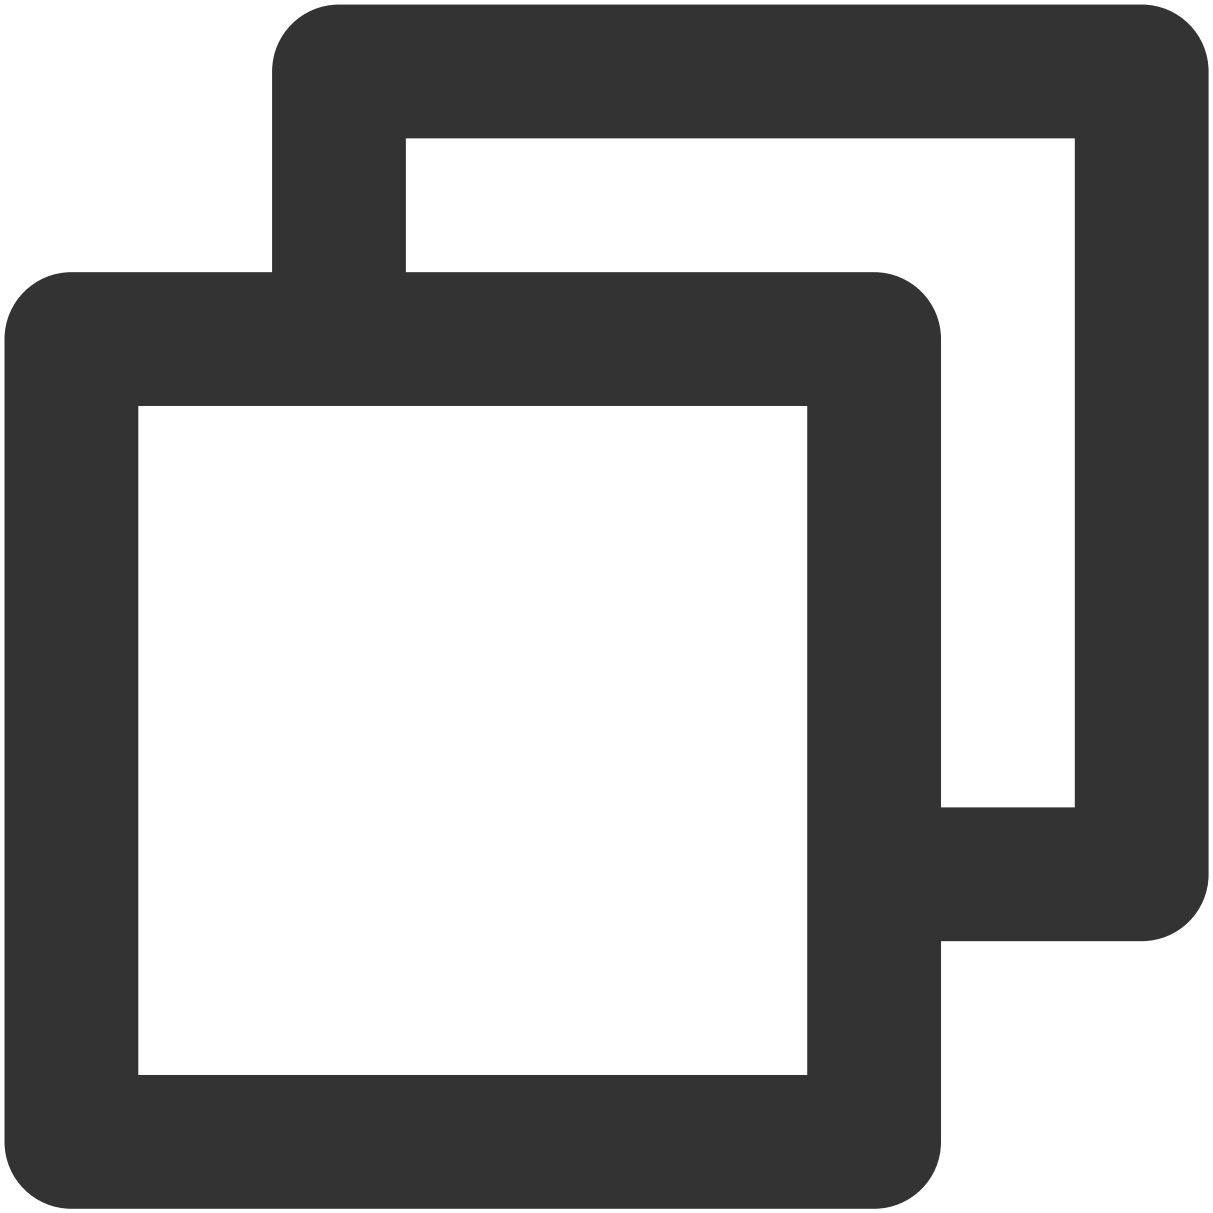

```
netstat -ano | findstr :80
```

For a Linux server, use the following command:

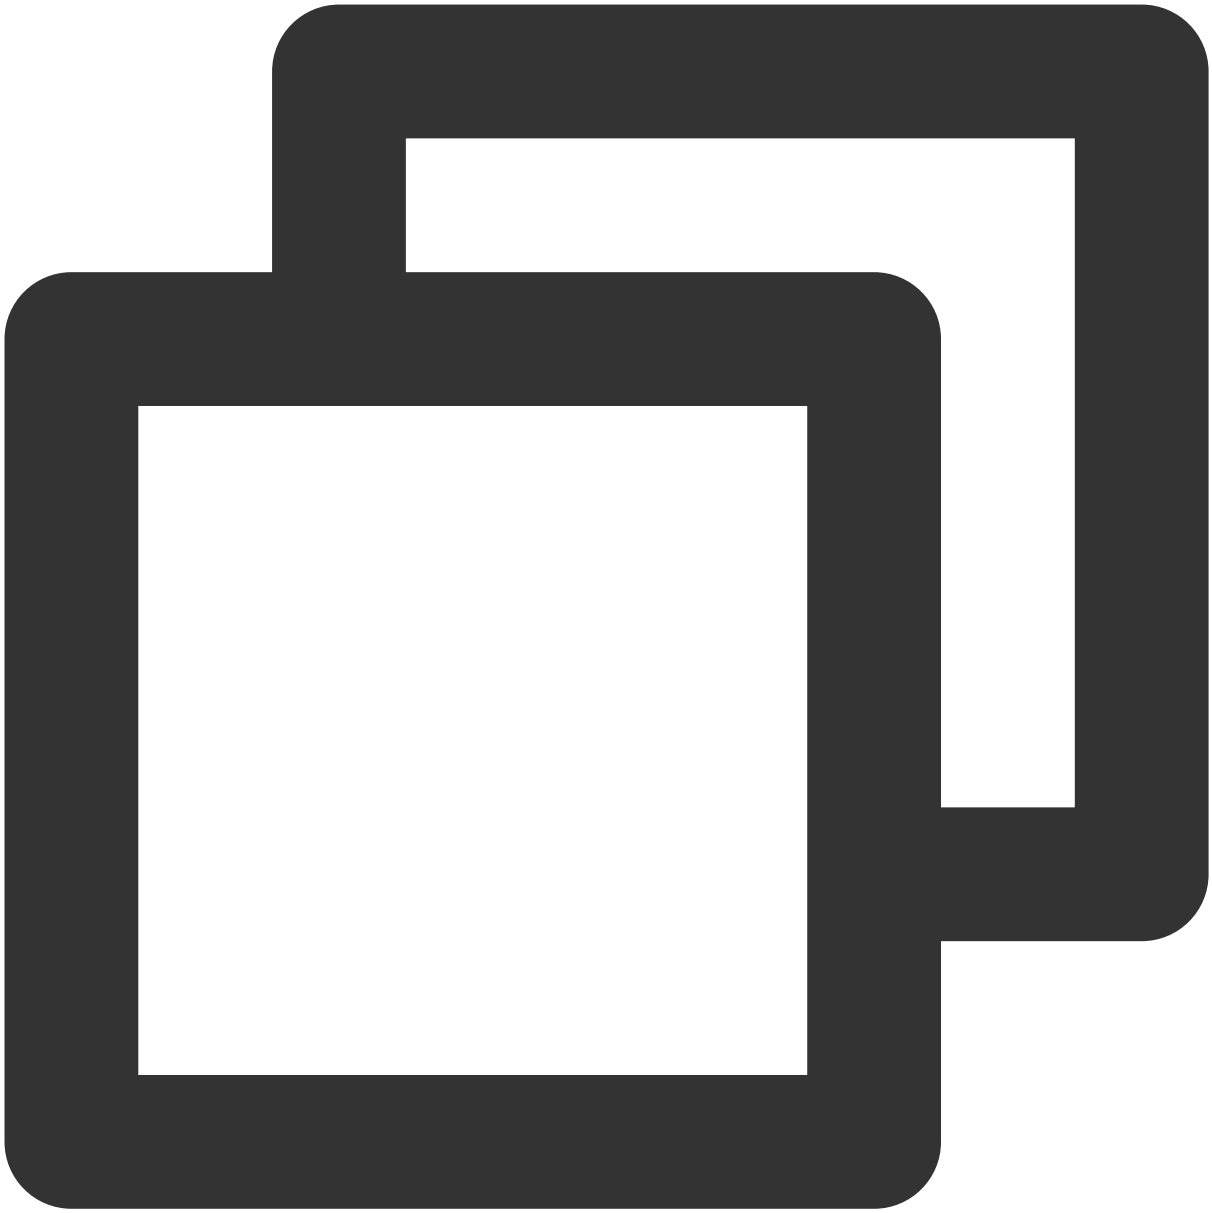

```
netstat -anp | grep :80
```

If you can see the listening of `1.1.1.10:80` or `0.0.0.0:80`, the configuration is correct.

Ensure that the backend port configured in CLB listener has been enabled on the real server.

For a layer-4 CLB instance, run the `telnet 1.1.1.10 80` command. If the backend port responds, the port is enabled.

For a layer-7 CLB instance, the backend port must respond with an HTTP status code that indicates success, such as status code 200. The check method varies depending on the operating system:

On Windows, enter the private IP address in the browser of a CVM instance to check whether it is normal. This example uses `http://1.1.1.10` .
On Linux, you can run the `curl -I` command to check whether the status is `HTTP/1.1 200 OK` . This example uses the `curl -I 1.1.1.10` command.
Check whether the backend CVM instance has a firewall or other security software, which may block the local IP address of the CLB instance. As a result, the CLB instance may fail to communicate with the real server.
Check whether the private network firewall of the server allows port 80 to pass. You can temporarily disable the firewall.
On Windows, run the `firewall.cpl' command in the Run command window to disable the firewall.
On Linux, run the `/etc/init.d/iptables stop` command to disable the firewall. If you use CenOS 7.x, run the `systemctl stop firewalld` command.
Check whether the health check parameters of CLB are configured correctly. We recommend you use the default health check parameter values in Health Check Overview.
For the test file specified for health check, we recommend you use a simple page in HTML format, which is used only to check the returned results. Dynamic programming languages such as PHP are not recommended.
Check whether the CVM instance has high load that leads to slow response.
Check the HTTP request method.
If the HEAD method is used, the real server must support HEAD.
If the GET method is used, the real server must support GET.
The health check result may be abnormal if both the `tcp_tw_recycle` and `tcp_timestamps` parameters are enabled. We recommend that you disable `tcp_tw_recycle` . For more information, see Solution to Excessive Clients in TIME_WAIT Status.

Highly Frequent Health Checks

For example, you configured to send a health check packet every 5 seconds in the console. However, the real server receives one or more health check requests per second. This is mainly related to the implementation mechanism of CLB health check.
Assume that 1 million requests from clients are distributed to four CLB servers, which then forward them to real servers. Each CLB server performs health checks separately. Therefore, if the CLB servers are configured to send a health check request every 5 seconds, each CLB server send a health check request every 5 seconds, and the real server may receive 4 health check requests within 5 seconds.

This implementation mechanism is efficient and accurate, and avoids false removal of servers. For example, if one of the eight physical servers in a CLB cluster fails, the other seven servers can still forward traffic normally.

If your business is load-sensitive, highly frequent health checks may cause business interruptions. You can reduce the business impact by setting a larger time interval, such as 15 seconds, for health checks.

If a real server is bound to multiple CLB instances, each CLB instance sends health check packets to detect the server health, resulting in a high frequency of health checks.

HTTPS

Last updated：2024-01-04 14:39:00

What cipher suites are supported by HTTPS?

ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES256-GCM-SHA384:DHE-RSA-AES128-GCM-SHA256:ECDHE-RSA-CHACHA20-POLY1305:DHE-RSA-CHACHA20-POLY1305:kEDH+AESGCM:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-AES256-SHA384:ECDHE-ECDSA-AES256-SHA384:ECDHE-RSA-AES256-SHA:ECDHE-ECDSA-AES256-SHA:AES128-GCM-SHA256:AES256-GCM-SHA384:AES128:AES256:AES:HIGH:!aNULL:!eNULL:!EXPORT:!DES:!RC4:!MD5:!PSK

[Back to Top]

What versions of SSL/TLS security protocols does HTTPS support?

The ssl_protocols supported by CLB HTTPS include TLSv1, TLSv1.1, TLSv1.2, and TLSv1.3.

[Back to Top]

What port can I use for HTTPS listening?

Not mandatory. Port 443 is recommended.

[Back to Top]

Why HTTPS mutual authentication is needed?

Some users such as financial service providers have higher requirements for data security. They require HTTPS authentication on both the server and client. To meet their needs, HTTPS two-way authentication is provided.

[Back to Top]

Why does the HTTPS actually generate more traffic than the billed traffic?

If the HTTPS protocol is used, it actually generates more traffic than the billed traffic as some of the traffic is used for protocol handshake.

[Back to Top]

Will requests from CLB instances to real servers still be transferred over HTTP after an HTTPS listener is added?

Yes. After an HTTPS listener is added, requests from a client to a CLB instance will be encrypted over HTTPS, but requests from a CLB instance to a real server will still be transferred over HTTP. Therefore, there is no need to configure SSL on real servers.

[Back to Top]

What types of certificates does CLB currently support?

CLB supports uploading the server certificate and CA certificate. For server certificate, the certificate content and private key need to be uploaded; for CA certificate, only the certificate content needs to be uploaded. Both certificates can be uploaded in PEM encoding format only.

[Back to Top]

How many HTTPS certificates can a listener be bound to?

If HTTPS one-way authentication is used, only one server certificate can be bound to a listener. If HTTPS mutual authentication is used, one server certificate and one CA certificate need to be bound to a listener.

[Back to Top]

How many cloud load balancers and listeners can one certificate be applied to?

A certificate can be applied to one or more cloud load balancers, or multiple listeners.

[Back to Top]

How do I upload a certificate?

You can upload it by calling an API or through the CLB console.

[Back to Top]

Is a certificate region-specific?

No. After the certificate is purchased and issued, its installation and deployment are not restricted by regions.

[Back to Top]

Do I need to upload the required certificates to real servers?

No. CLB HTTPS provides a certificate management system to manage and store user certificates. Certificates do not need to be uploaded to backend CVM instances, and all the private keys uploaded to the certificate management system are stored in an encrypted manner.

[Back to Top]

What should I do after the certificate expires?

You need to manually update the certificate.

[Back to Top]

What can I do when a certificate error occurs?

The error may occur due to incorrect private key. You need to replace the certificate with a new one that meets business requirements.

[Back to Top]

WS/WSS Protocol Support

Last updated：2024-01-04 14:39:00

Product In
troduction

What is WS/WSS?

WebSocket (WS) is a protocol that provides full-duplex communication channels over a single TCP connection. WebSocket facilitates data exchange between the client and server, and allows active data push from the server to client. In WebSocket API, only one handshake is required between the browser and server to create a persistent connection and carry out bi-directional data transmission.

[Back to Top]

Why should WS/WSS be used?

Without WebSocket, the client has to pull data from the server through polling.
There are two shortcomings in this data exchange method:
1. Low efficiency. To pull real-time data, the client has to frequently initiate the Ajax request.
2. The server cannot push data proactively.
WebSocket is designed to solve these problems. As a new protocol released when HTML5 was launched, WebSocket achieves full-duplex communication between the browser and server. It can transmit message-based text and binary data, solving HTTP problems at the protocol level.
Key advantages of WebSocket:
1. Less overhead. After the connection is established, the packet header used for control is small. Compared to an HTTP request that requires a complete header, WebSocket helps reduce the overhead.
2. Real-time push. As a full-duplex protocol, WebSocket can achieve real-time data push from server to client.
3. Persistent connection.

[Back to Top]

How is WS/WSS billed?

CLB supports WS/WSS by default and no additional fees will be charged.

[Back to Top]

How do I enable WS/WSS for CLB?

WS/WSS is enabled for CLB by default. If a connection is idle for more than 60s, you need to customize the `proxy_read_timeout` parameter, which should be less than 900s preferably. For more information, see Layer-7 Custom Configuration.

If the listener listens to HTTP or TCP SSL, WS is supported by default. If it listens to HTTPS, WSS is supported by default.

When WSS is used, CLB will carry out SSL offloading.

[Back to Top]

Which regions support WS/WSS?

Currently, WS/WSS protocols are supported in **all regions**.

[Back to Top]

HTTP/2 Protocol Support

Last updated：2024-01-04 14:39:00

What is HTTP/2?

HTTP/2 (Hypertext Transfer Protocol Version 2) is a major revision of the HTTP network protocol used by the World Wide Web.
HTTP/2 is designed to address the performance issues in HTTP1.X to better use network resources and reduce network application latency.
HTTP/2 is backward compatible with HTTP1.X.
[Back to Top]

Why should I use HTTP/2?

Compared with HTTP1.X, HTTP/2 can make the response be more fast and efficient. HTTP/2 has the following advantages:
Multiplex: concurrent processing brings a faster response.
Server push: the server proactively pushes resources needed by the client, reducing the number of requests.
More features include bandwidth limit, request priority, header compression, and binary framing.
[Back to Top]

How does the billing work?

CLB supports the HTTP/2 protocol without charging extra fees.
[Back to Top]

How do I enable HTTP/2 on CLB?

Note：
The HTTP listener does not support HTTP/2. Mainstream browsers and web servers only support the TLS-based HTTP/2 protocol.
The HTTP1.X protocol is still used between the CLB instance and the real server.
1. Enable HTTP/2 on HTTPS listeners
CLB instance: you can enable or disable the HTTP/2 protocol in a CLB instance. For more information, please see Configuring an HTTPS Listener.
Classic CLB instance: HTTPS listeners created for a classic CLB instance before April 2018 do not support HTTP/2. HTTPS listeners created after April 2018 support but cannot disable HTTP/2.
2. Agree on the protocol at client access
When the client accesses an HTTP/2-enabled listener, the protocol version will be negotiated during the handshake process of HTTPS. The client uses ALPN (Application-Layer Protocol Negotiation) to inform the server of a list of supported protocols. The server selects HTTP/2 or HTTP1.X according to the protocol list. If the client does not support HTTP/2, the server will be automatically backward compatible without requiring additional configuration.
[Back to Top]

Which regions support HTTP/2?

Currently, all regions support HTTP/2.
[Back to Top]

Default Domain Name Blocking Prompt

Last updated：2024-04-16 14:46:20

According to the relevant national laws and regulations, After then, you cannot use the default domain name to access Tencent Cloud. For example, if `xxxxxxx-gz-tencentclb.com` is detected for the Guangzhou region, access to it will be denied.
On the platform side, it is recommended to resolve your custom domain name to the default domain name provided by the platform side using CNAME. For operation details, see Configure CLB Forwarding Domain Name.

Advantages of Configuring Your Custom Domain Name

Enhance Brand Image: Configure a custom domain name to your service. Your personalized domain name enhances brand image and professionalism and increases user trust.
Prevent Domain Name Blocking: Some applications or platforms may block the default domain name of Cloud Load Balancer. By binding your custom domain name, you can ensure that your service is always accessible.
Enhance Access Experience: Accessing services with your custom domain name makes it easy for your users to memorize. Compared with using a default domain name, it is more concise and user-friendly, allowing easy access and sharing.
Ensure Service Continuity: After binding your custom domain name to the service, users can still access your service using the same domain name even if the services change later, ensuring link persistence and long-term accessibility.