Cloud Load Balance

FAQ

Product Introduction

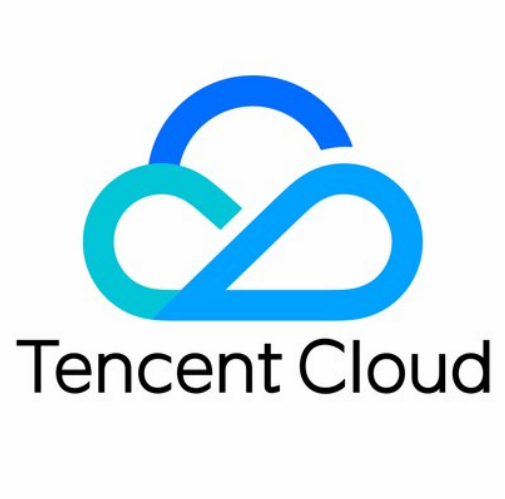

Contents

FAQ
FAQ for Load Balancing Configurations

Last updated：2018-06-13 09:59:47

1. What to do with a CVM instance exception indicated in the health check?

Please conduct troubleshooting using the following steps:

- Make sure that you access your application service directly via the CVM.
- Make sure that the relevant port is enabled on the backend CVM.
- Check whether there is security software like firewall inside the backend CVM because it can make the cloud load balancer systems fail to communicate with the backend CVM.
- Check if the parameter settings of cloud load balancer are correct.
- It is recommended to use static pages for health check.
- Check whether there is a high load on the backend CVM that leads to a slow response of the CVM.
- Make sure that there are no iptables restrictions on the submachine of CVM.

2. What to do if no policy file is returned and the connection is directly broken after a policy request (i.e. flash server request) is sent from port 843?

When the cloud load balancer has received policy request from port 843, it will return the common crossdomain policy configuration file. If no policy file is returned and the connection is directly broken, this may be caused by the incorrect flash server request.

Make sure to send the flash server request correctly: \0.

> Note: It must be ended with \0, and a total of 23 bytes are allowed. \0 stands for a character with accii code 0 and only occupies one byte.

Normally, the returned result of 843 request is as follows:

```
VM_02_sles10_64:/ # perl -e 'printf "<policy-file-request/>%c",0' | netcat -i 1 101.226.62.63 843
<?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>
```

3. Is it supportive to acquire the real IP of client?

Tencent Cloud's IP acquisition capability is automatically enabled, and the real IP of client can be acquired in the X-forwarded-for mode.

4. Which TCP ports can cloud load balance be applied to?

You can perform cloud load balance for the following TCP ports: 21 (FTP), 25 (SMTP), 80 (Http), 443 (Https), 1024-65535, etc.

5. How does cloud load balancer achieve session persistence based on cookies?

In the Cookie Insertion mode, CLB is responsible for inserting cookies without making any modification to the backend CVM. When the client makes the first request, the client HTTP request (without cookie) goes into CLB, which then selects a backend CVM based on the cloud load balance algorithm policy and sends the request to the CVM. Then the backend CVM gives an HTTP reply (without cookie), which is sent back to CLB, and then CLB inserts the cookie into it and returns the HTTP reply (with cookie) to the client.

When the client makes the second request, the client HTTP request containing the cookie inserted by CLB last time goes into CLB, which then reads the session persistence values in the cookie and sends the HTTP request (with the same cookie as above) to the specified CVM. Then the back-end CVM gives a reply, and because the CVM does not write the cookie, the HTTP reply does not contain the cookie. When the reply traffic flows into the CLB again, the updated session persistence cookie will be written to CLB.

6. What is the difference between Layer-4 and Layer-7 cloud load balancers?

Layer-4 cloud load balance capability is based on IP and port, while Layer-7 cloud load balance capability is based on the application layer information such as HTTP header and URL.

The difference between Layer-4 and Layer-7 cloud load balancers lies in whether Layer-4 information or Layer-7 information is used as the basis for determining the way of forwarding traffic when cloud load balance is performed on backend CVMs. For example, Layer-4 cloud load balancer determines which traffic needs load balance based on Layer-3 IP address (VIP) and Layer-4 port number, performs NAT on the traffic to be processed and then forwards it to the backend CVM. At the same time, it records which CVM has processed the TCP or UDP traffic, and forwards all the subsequent traffic of this connection to the same CVM for processing.

Layer-7 cloud load balancer takes application layer's characteristics into account on the basis of Layer-4 cloud load balancer. For example, for the same Web server, in addition to determining the traffic to be processed based on VIP and Port 80, Layer-7 cloud load balancer can decide on whether to perform load balance based on URL, browser category and language at Layer 7. Layer-7 cloud load balancer is also known as "content exchange", which is designed to decide on the final choice of internal CVM based on the really meaningful application layer content in message and CVM selection method set for the cloud load balancer.

To select CVM based on the real application layer content, Layer-7 cloud load balancer must establish a connection with the client as a proxy of the final CVM (three-way handshake) to receive the message containing real application layer content from client, and then determines the final choice of the internal CVM according to the specific fields in the message plus the server selection method set for the cloud load balancer. In this case, the cloud load balancer is more like a proxy server. The cloud load balancer establishes TCP connection with frontend client and backend CVM separately.

7. Can CVM forward traffic from port A to another port on the same server by configuring private network cloud load balancer?

No. For the access to port a on server A (10.66.*.101), the request can be forwarded to port b on server B (10.66.*.102) through private network cloud load balancer, but the traffic cannot be forwarded to port b on the same server A (10.66.*.101).

8. What is the backend CVM weight?

Users can specify the forwarding weight for each CVM in the backend CVM pool, and the CVM with a higher weight ratio will be assigned more access requests. Users can set the weight for backend CVMs individually based on their service capabilities and statuses.

If you have enabled session persistence at the same time, the accesses to the backend application servers may be not exactly the same. You're recommended to temporarily disable the session persistence and then observe whether the problem still exists.

9. What is the difference between UDP and TCP?

TCP is a connection-oriented protocol. Therefore, when TCP is used, it is necessary to establish a reliable connection with the other side before receiving and sending data. UDP is a non-connection-oriented protocol, and it directly sends data packets without performing three-way handshake with the other side. UDP is suitable for the scenarios focusing more on real-timeness than reliability, such as video chat, real-time push of financial market information, DNS, Internet of Things.

10. Does the backend CVM need public network bandwidth? Will it affect the service of cloud load balancer?

No traffic or bandwidth fee is charged for cloud load balancer. Any public network traffic generated by the cloud load balance service will be charged to the bill for the backend CVM. You're recommended to choose "Bill-by-Traffic" for public network bandwidth when purchasing backend CVM and setting a reasonable peak bandwidth cap, so that you do not need to keep track of the fluctuation of total traffic of CLB egresses. Web traffic on the Internet has a considerable fluctuation that cannot be predicted accurately. When "Bill-by-Bandwidth" is selected, it is not cost-effective if excessive bandwidth is purchased, and packet loss may occur during business peak if insufficient bandwidth is purchased.

11. HTTP redirection in load forwarding

When you visit the website `http://example.com` through a browser, a redirection to the root directory is required for the CVM. When you visit the website `http://example.com/` through the browser, the CVM will directly return the default page of root directory set by the site. Similarly, if `http://cloud.tencent.com/movie` is redirected to `http://cloud.tencent.com/movie/` through URL rewriting, then 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 "/" behind the secondary URL.

In Tencent Cloud's cloud load balancer, if frontend port number is different from the backend one, "/" needs to be added behind the secondary URL upon the visit to secondary page to avoid the change of port number after the HTTP redirection and ensure the normal visit to the page.

Assume that in Layer-7 forwarding, port 80 on the CLB instance and port 8081 on the backend CVM are listened. If the client accesses `http://www.example.com/movie` , the access request is forwarded to the backend CVM via cloud load balancer, and then the CVM redirects the request to `http://www.example.com:8081/movie/` (listened port is 8081). In this case, the client access fails (port error).

Therefore, it is recommended to rewrite the access request to the secondary URL with "/", such as `http://www.example.com/movie/` . This can avoid HTTP redirection, eliminate the need to make unnecessary judgment and reduce unwanted load. If it is necessary to use HTTP redirection, make sure that the cloud load balancer's listened port is the same as that of the backend CVM.

12. Notes about compatible versions in case of inconsistency of HTTP versions between client and server

Forwarding Compatibility

- For the frontend (client), HTTP1.0/1.1 and backward compatibility are supported.
- For the backend (server), Tencent Cloud uses HTTP1.0 protocol. HTTP1.0/1.1 and backward compatibility are supported.

> Note: HTTP/2 is only supported in HTTPS, and backward compatibility is allowed on both client and server. HTTP protocol is not supported.

Support for Gzip Compatibility

- For the frontend (client), HTTP1.0/1.1 and backward compatibility are supported. Additional configuration is not needed, since mainstream browsers all support Gzip.
- For the backend (server), as HTTP/1.1 protocol is supported on the CVM over Tencent Cloud private network, you don't need to make any configuration. HTTP/1.1 is configured by default using nginx, thus achieving compatibility.

Note: HTTP/2 is only supported in HTTPs, but Gzip can be used in any HTTP versions supported by Tencent Cloud.

13. How to make the settings for the security group of cloud load balancer backend CVMs? How to set the access blacklist?

Security Group Configuration for Cloud Load Balancer

If security group rules have been set for the backend CVM, it is likely that the cloud load balancer instance cannot communicate with the CVM. Therefore, in Layer-4 forwarding and Layer-7 forwarding, it is recommended to set the security group of the backend CVM to "Allow ALL". If the security group is enabled and deny access from any protocols and IP segments by default, IPs of all clients need to be configured to the security group rules of the IP of the CVM.
For some malicious IPs, you can add these IPs to the top rules of the security group to prevent them from accessing the backend CVM; and then allow the accesses from all IPs (0.0.0.0) to the local service port, so that the normal clients can access the CVM. Security group rules are arranged in priority order and are matched from top to bottom.

If health check has been set for Layer-7 cloud load balancer forwarding in VPC, please also note that the cloud load balancer's VIP needs to be added to the allow rule of the backend CVM security group, otherwise the health check may fail.

Setting Access Blacklist

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

- Add the client IPs from which access needs to be rejected plus ports to the security group, and choose "reject the access from the IPs" in the Policy section.
- When the setting is made, add another security group rule that allows all accesses to the port from all IPs by default.
  When the configuration is completed, the security group rules are as follows:

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

Note:***The above configuration steps should be performed in a correct order, otherwise the blacklist configuration cannot take effect.***

For more information about security groups, please see Access Control for the Backend CVM

14. Notes about the ability of private network CLB to directly acquire client IP

The following notes apply to the private network CLB (choose private network VPC) purchased after October 24, 2016. The private network-based CLB no longer performs SNAT processing. The access IP acquired from the server is the real client IP and no additional configuration is required. To ensure that your business operates properly, note the followings:

1. For the private network-based CLB purchased after October 24, 2016, when the security group policy is enabled, the inbound access from client IP must be allowed to ensure normal access.

2. If necessary, you can switch the existing private network-based CLB to the new one by submitting a ticket to the after-sales team. After the switching, the access IP acquired from the server side is the client IP. A several-minute business interruption may occur during the switching.

15. Notes about too frequent health check

Health check packets are sent too frequently. Each health check packet is sent every 5 seconds as configured in the console. But the backend RS finds that one or more health check requests are received in one second. Why is that?

Too frequent health check is caused by the implementation mechanism of CLB backend health check. Assume that 1 million requests from client are distributed on four LB backend physical machines before being sent to the CVM, and each LB backend physical machine conducts health check separately. If the LB instance is set to send a health check request every 5 seconds, each physical machine on the LB backend sends a health check request every 5 seconds. That's why the backend CVM receives multiple health check requests. For example, if the cluster to which an LB instance belongs has eight physical machines, and each machine sends a request every 5 seconds, the backend CVM may receive 8 health check requests in 5 seconds.

The advantages of this implementation solution are high efficiency, accurate check, and avoidance of mis-removal. For example, if one of eight physical machines in the LB instance cluster fails, the other seven machines can still forward traffic normally.

Therefore, if your backend CVM is checked too frequently, you can set the check interval to be much longer, such as, 15 seconds.

FAQ for HTTPS

Last updated：2018-03-28 17:21:41

1. Which encryption 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

2. Which versions of SSL/TLS security protocols are supported by HTTPS?

Cloud Load Balancer HTTPS currently supports the following SSL protocols: TLSv1, TLSv1.1, TLSv1.2

3. Which port should be used for HTTPS listener?

There is no mandatory requirement for this. Port 443 is recommended.

4. Why is HTTPS mutual authentication needed?

Some customers have a higher requirement for data security, such as those who deal with financial services. They require HTTPS authentication to be carried out on both server and client. In order to cater for the needs of such customers, we have launched HTTPS mutual authentication.

5. Why does HTTPS protocol actually generate more traffic than the billed traffic?

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

6. When the HTTPS listener has been added, are the requests between the Cloud Load Balancer and back-end CVM still transmitted over HTTP protocol?

Yes. When the HTTPS Listener has been added, requests between the client and Cloud Load Balancer are encrypted over HTTPS protocol, and requests between the Cloud Load Balancer and back-end CVM are still transmitted over HTTP protocol. Therefore, there is no need to make SSL configuration for back-end CVM.

7. What types of certificates are supported by CLB currently?

It currently supports the upload of server certificate and CA certificate. For server certificate, both certificate content and private key are required to be uploaded. For CA certificate, only certificate content is required to be uploaded. Both types of certificates only support upload in PEM encoding format.

8. How many HTTPS certificates can be bound to a listener?

If HTTPS unidirectional 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 are required to be bound to a listener.

9. How many Cloud Load Balancers or listeners can a certificate be applied to?

A certificate can be applied to one or multiple Cloud Load Balancers, or multiple listeners.

10. How to upload a certificate?

You can upload it by using API or Cloud Load Balancer console.

11. Are the certificates region-sensitive?

Yes. If a user's certificate needs to be used in multiple regions, it is necessary to upload the certificate in multiple regions to ensure the security and performance.

12. Can a certificate be deleted after being uploaded?

The Deletion function is not available now.

13. Do the certificates need to be uploaded to back-end CVM?

No. Cloud Load Balancer HTTPS provides the Certificate Management System to manage and store user certificates. Certificates do not need to be uploaded to back-end CVM, and all the private keys uploaded to the Certificate Management System are stored in an encrypted form.

14. What to do after a certificate expires?

When the current certificate expires, users need to update the certificate manually.

15. What to do when an error occurs while adding a certificate?

This maybe caused by the wrong content of private key. In this case, user need to replace it with a new certificate that meets the requirements.

WS/WSS Protocol Support

Last updated：2018-06-01 17:12:02

Product Overview

What is WS/WSS?

WebSocket is a protocol which performs full-duplex communication over the single TCP connection. WebSocket makes it easier to exchange data between client and server, and allows active data push from server to client. In WebSocket API, only one handshake between browser and server is required before a persistent connection between them is created and two-way data transmission is carried out.

Why do we use WS/WSS?

Without WebSocket, the client had to pull data from the server through polling when data on the server are needed.
There are two shortcomings in this data exchanging method:

1. Low efficiency. When real-time data are needed, the client has to initiate Ajax requests frequently to pull data.
2. The server cannot push data actively.
   WebSocket is designed to solve these problems. WebSocket is a new protocol released with HTML5. It achieves full-duplex communication between browser and server, and can transmit message-based texts and binary data. WebSocket solves the above problems of HTTP at protocol level.

Key advantages of WebSocket:

1. Lower control cost After the connection is established, the header used for control is small. A complete header is required in each HTTP request, but there is no such requirement in WebSocket, so the corresponding cost can be reduced significantly.
2. Satisfactory real-timeness. As a full-duplex protocol, WebSocket can achieve the real-time data push from server to client.
3. Capable of maintaining connection status.

Product Purchase

How is WS/WSS billed?

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

Product Implementation

How to enable WS/WSS on CLB?

WS/WSS is enabled by default and no additional configuration is required.
If the listener listens to HTTP, WS is supported by default. If it listens to HTTPS, WSS is supported by default.
When WSS is used, CLB will unmount SSL.

Which regions support WS/WSS?

WS/WSS protocols are supported for **All regions**.