ALIBABA CLOUD

Alibaba Cloud

Server Load Balancer

FAQ

Document Version: 20220322

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	**Danger:** Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	**Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	**Notice:** If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	**Note:** You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.Why am I unable to access an SLB instance?

This topic describes why a Server Load Balancer (SLB) instance fails to be accessed from a client and how to troubleshoot the issue.

> **Note** In this example, the frontend port number of the SLB instance is 80, the port number of the backend ECS instance is 80, and the internal IP address of the ECS instance is 10.11.192.1. You must configure the ports and internal IP address based on your business requirements.

No.	Possible cause	Solution
1	SLB cannot be accessed by backend servers. For Layer 4 SLB, an ECS instance cannot provide services for clients and function as the backend server of the SLB service at the same time.	N/A
2	A health check exception occurs.	For more information about how to troubleshoot health check exceptions, see How do I troubleshoot health check exceptions of a layer-4 (TCP/UDP) listener? and How do I troubleshoot a health check exception of a layer-7 (HTTP/HTTPS) listener?.
3	SLB does not support FTP, TFTP, H.323, and SIP protocols.	• For a Linux system, you can configure the forwarding of port 22 and use SFTP to connect and transmit data. • You can associate an elastic IP address (EIP) with an FTP server in cut-through mode to provide external FTP service. For more information, see Deploy an FTP server by using an EIP.

4	The internal firewall of the server does not allow traffic on port 80.	You can run the following commands to temporarily disable the firewall to perform a test.
		• For a Windows server, run the following command:
		`firewall.cpl`
		• For a Linux server, run the following command:
		`/etc/init.d/iptables stop`

No.	Possible cause	Solution
5	A backend port exception occurs.	• For Layer 4 SLB, you can perform a Telnet test. If you receive a response, the backend port functions properly. For example, you can run the following command to perform a test: `telnet 10.11.192.1 80`. • For Layer 7 SLB, you can check the returned HTTP status code. The status code must be a status code that indicates a normal condition, such as 200. You can use the following methods to test whether the backend port is normal: ◦ Windows: Access the internal IP address of the ECS instance to check the connectivity. Example: `http://10.11.192.1` ◦ Linux: Run the `curl -I` command to check whether the status is `HTTP/1.1 200 OK`. Example: `curl -I 10.11.192.1`

6	The rp_filter feature conflicts with the policy-based routing mechanism of the Linux Virtual Server (LVS) for SLB.	1. Log on to the ECS instance that is attached to Layer 4 SLB. The ECS instance runs a Linux system.
		2. Edit the */etc/sysctl.conf* file and set the following parameters in the system configuration file to 0:
		`net.ipv4.conf.default.rp_filter = 0`
		`net.ipv4.conf.all.rp_filter = 0`
		`net.ipv4.conf.eth0.rp_filter = 0`
		3. Run the `sysctl -p` command to make the configurations take effect.

7	A listener exception occurs.	Run the following commands on the server. If you can see the listening information of 10.11.192.1:80 or 0.0.0.0: 80, the listening on the port is normal.
		• For a Windows server, run the following command:
		`netstat -ano | findstr :80`
		• For a Linux server, run the following command:
		`netstat -anp | grep :80`

8	No listeners are configured for the SLB instance.	Configure one or more listeners. For more information, see Listener overview.
9	The SLB instance cannot be accessed by using its domain name. This may be caused by an error in domain name resolution.	N/A

No.	Possible cause	Solution
10	An exception occurs on the on-premises network of the client or the intermediate link of the service provider.	Test the connectivity on the service port of the SLB instance in different regions and network environments. If the exception occurs only when the SLB instance is accessed from the on-premises network, it can be determined that the problem is caused by a network exception. You can perform ping and MTR tests for further troubleshooting and analysis.
11	The client IP address is blocked by Alibaba Cloud Security.	1. Visit `http://ip.taobao.com` in the client network environment to obtain the public IP address of the client. 2. Add the IP address to the SLB whitelist to allow access from the IP address. **Note** This operation may pose security risks. Make sure that the IP addresses in the whitelist do not incur malicious attacks on SLB.
12	After you switch from Anti-DDoS Pro or Anti-DDoS Premium to Anti-DDoS Basic, the whitelist is not disabled.	Disable the whitelist.
If your problem persists, submit a ticket and provide the following information: • The ID of the SLB instance or the IP address of the SLB instance • The public IP address of the client obtained when you visit `ip.taobao.com` • Screenshots of the client running ping and MTR tests by using the IP address of the SLB instance		

2.Why is the traffic among my ECS instances unevenly distributed?

Causes

Traffic may be unevenly distributed due to the following reasons:

- Only a small number of requests are being received by ECS instances.
- The target ECS instances have different network capacities.

  > **Note** The memory usage of ECS instances does not indicate whether requests are evenly distributed.

- Session persistence is enabled.

  If session persistence is enabled, it will cause traffic imbalance when few clients are accessing the Server Load Balancer (SLB) instance. This is especially common when a small number of clients are used to test the SLB instance. For example, session persistence (based on source IP addresses) is enabled for a TCP listener and a client is used to test the load balancing service.

- The ECS instance status is abnormal.

  Backend servers with abnormal heath status can also lead to an imbalance especially during a stress test. If the health check for a backend ECS instance fails or the health status of a backend ECS instance changes frequently, this will cause an imbalance.

- TCP Keepalive is enabled.

  When some backend ECS instances enable TCP Keepalive and others do not, the connections will accumulate on the ECS instances with TCP Keepalive enabled. This scenario will cause an imbalance.

Troubleshooting

- Check whether the weights of backend ECS instances are the same.
- Check whether health checks of backend ECS instances fail or whether the health status is unstable in a specified period. Check whether the health check is correctly configured with the status code.
- Check whether both the WLC scheduling algorithm and session persistence are enabled. If so, change the scheduling algorithm to WRR.

3.What can I do if health checks generate an excessive number of logs?

Server Load Balancer (SLB) can automatically save health check logs generated in three days. If too many health check logs are generated and affect your maintenance, you can reduce health check logs or prevent certain logs from being generated through the following methods.

> **Note** If you reduce health check logs, SLB faults may be missed. Therefore, we recommend that you consider the risks of the following methods and use the methods with caution.

- Get access logs
- Adjust health check frequency
- Close Layer-7 health checks
- Change Layer-7 SLB to Layer-4 SLB
- Disable application logs on the health check page

Get access logs

HTTP health checks use the HEAD request method by default. Therefore, you can obtain access logs by filtering out HEAD requests.

Adjust health check frequency

You can increase the interval between two health checks to reduce the health check frequency and generated logs.

Potential risks

After you increase the interval, if a backend ECS instance fails, the time needed for SLB to detect the faulty ECS instance is increased accordingly.

Procedure

1. Log on to the SLB console.
2. On the **Server Load Balancer** page, click the ID of the target SLB instance.
3. On the **Listeners** tab, find the target listener, and click **Configure** in the **Actions** column.
4. On the **Configure Listener** page, click **Next** and then click **Next** again to go to the **Health Check** tab.
5. Adjust the **Health Check Interval**. Value range: 1 to 50. Unit: seconds. The greater the interval is, the lower the health check frequency is, and the fewer logs are generated by backend servers. Modify the interval according to your actual situation.

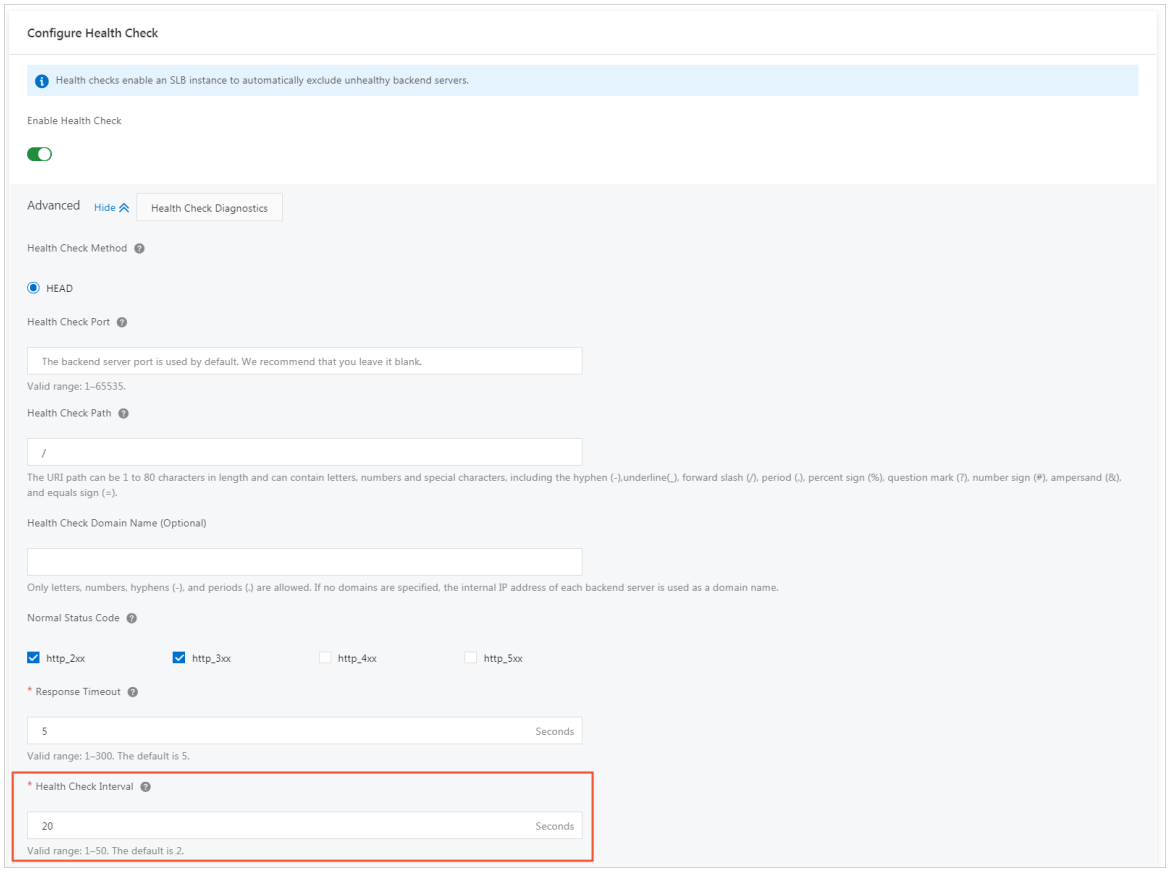

6. Click **Next** until the **Submit** tab is displayed. Then, click **OK**.

Close layer-7 health checks

When layer-7 (HTTP or HTTPS) SLB is used, health checks are performed through HTTP HEAD requests. Application logs of backend servers record the health check requests, leading to a large number of logs.

Potential risks

After you close HTTP/HTTPS health checks, SLB does not check backend servers. If a backend server fails, the traffic cannot be automatically forwarded to other normal backend servers.

Procedure

1. Log on to the SLB console.
2. On the **Server Load Balancer** page, click the ID of the target SLB instance.
3. Click the **Listeners** tab, find the target listener, and click **Configure** in the **Actions** column.
4. On the **Configure Listener** page, click **Next** and then click **Next** again to go to the **Health Check** tab.
5. Turn off **Enable Health Check**.

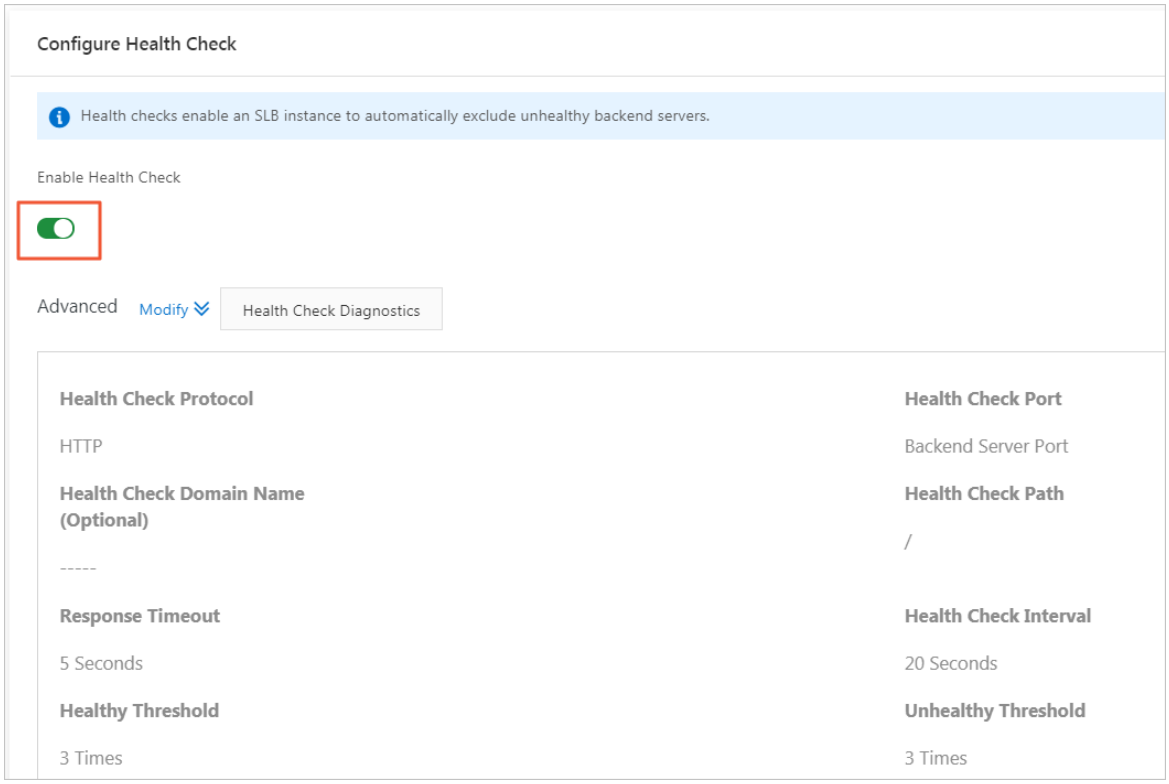

6. Click **Next** until the **Submit** tab is displayed. Then, click **OK**.

Change layer-7 SLB to layer-4 SLB

Layer-4 health checks are preformed through TCP three-way handshakes and generate no application logs. If you change layer-7 SLB to layer-4 SLB, the number of application logs can be reduced.

Potential risks

After you change the layer-7 SLB to layer-4 SLB, SLB checks only the status of the listener port and does not check the HTTP status. In this way, SLB cannot detect the exceptions occurring to HTTP applications in real time.

Procedure

1. Log on to the SLB console.
2. On the **Server Load Balancer** page, click the ID of the target SLB instance.
3. On the **Listeners** tab, find the target listener, and click **Configure** in the Actions column.
4. On the **Configure Listener** page, click **Next** and then click **Next** again to go to the **Health Check** tab.
5. Change the **Health Check Protocol** to **TCP**.

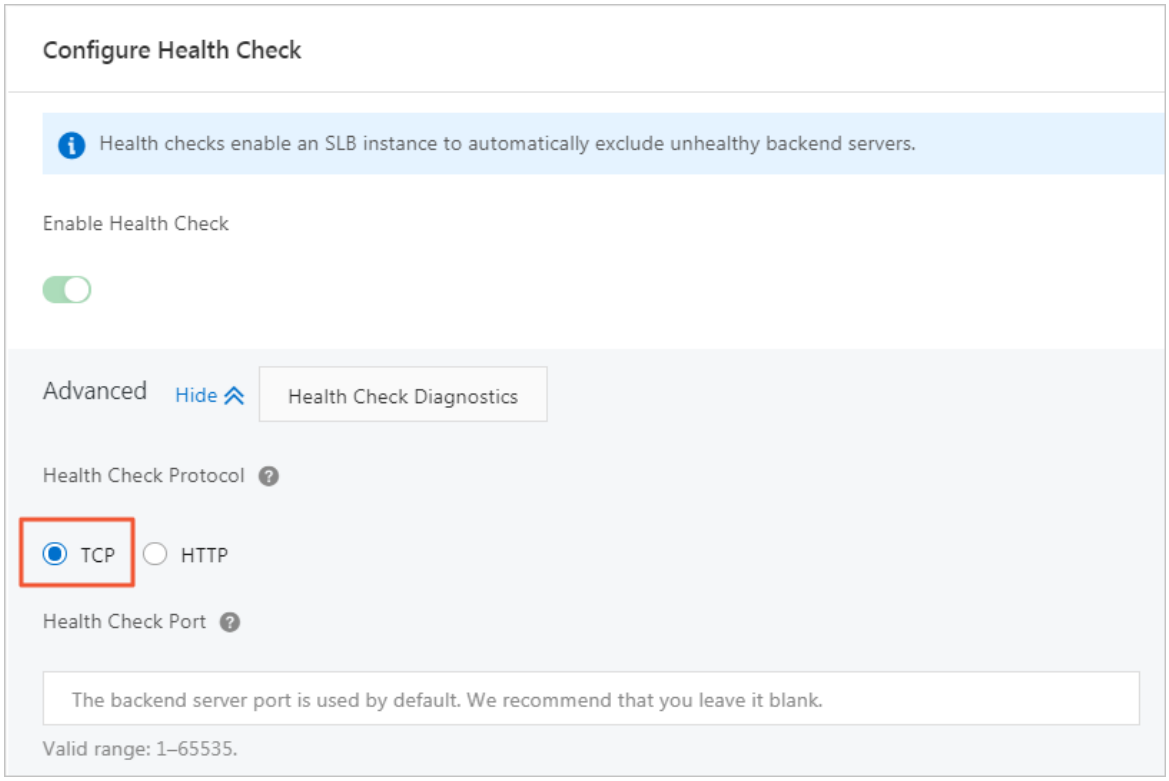

6. Click **Next** until the **Submit** tab is displayed. Then, click **OK**.

Disable application logs on the health check page

You can configure an independent site for health checks and disable application logs of this site. This method can also reduce the number of health checks. For example, the service site is abc.123.com. You can use test.123.com as the health check site and disable logs of test.123.com.

Potential risks

If the health check site is running normally, but an exception occurs to the service site, health checks cannot detect the exception of the service site.

Procedure

1. Create a new health check site and health check page on the backend server and disable logs. In this example, NGINX is used.

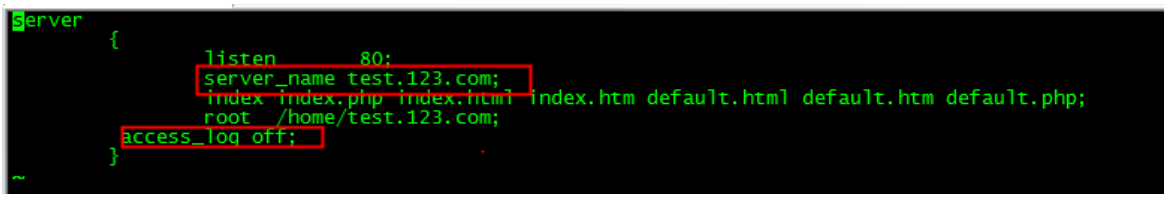

2. Log on to the SLB console.
3. On the **Server Load Balancer** page, click the ID of the target SLB instance.
4. On the **Listeners** tab, find the target listener, and click **Configure** in the **Actions** column.
5. On the **Configure Listener** tab, click **Next** and then click **Next** again to go to the **Health Check** tab.
6. In the **Health Check Domain Name (Optional)** field, enter the domain name of the health check site. In the **Health Check Path** field, enter the path of the health check page.

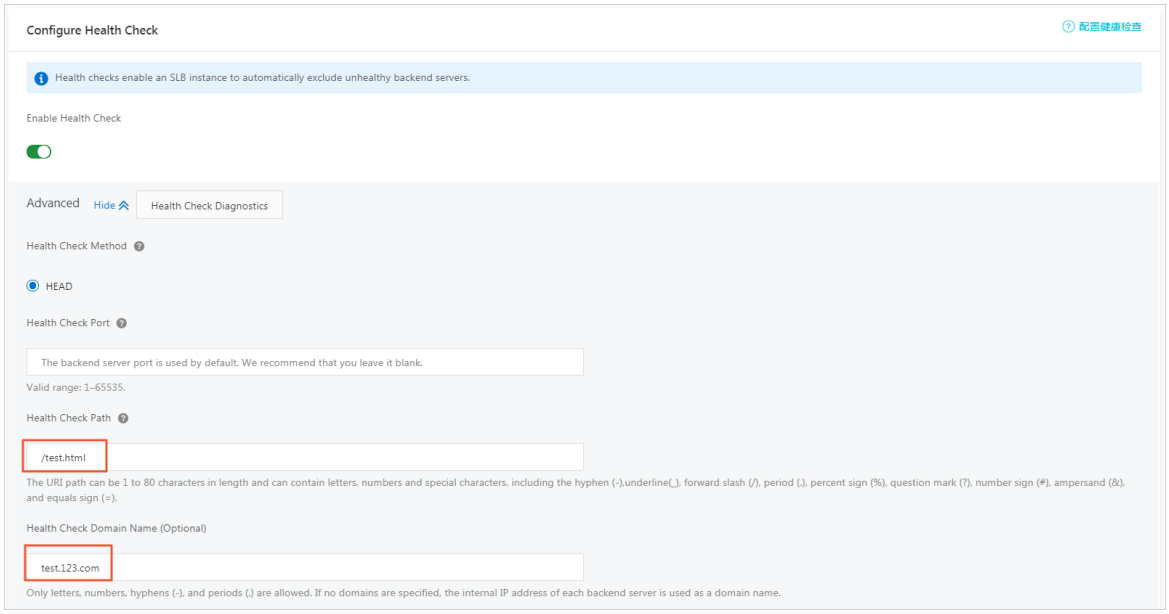

7. Click **Next** until the **Submit** tab is displayed. Then, click **OK**.

4.What can I do if my ECS instance is declared unhealthy after I enable health checks for Server Load Balancer?

After you enable health checks of Server Load Balancer (SLB), when one backend ECS instance is declared as unhealthy, requests are forwarded to other normal ECS instances. When the faulty ECS instance becomes normal, SLB forwards requests to the ECS instance again.

For layer-7 SLB service, when an ECS instance is declared as unhealthy, you can troubleshoot problems from the following aspects:

- Make sure you can directly access your service through the ECS instance.
- Make sure the backend port you configured in the listener is opened on the backend server.
- Check whether the backend ECS instance has installed a firewall or other security protection software. This type of software may block the local IP address of the SLB service, and thus disable the communication between the SLB service and the backend server.
- Check whether the SLB health check parameters are correctly set. We recommend that you use default health check settings.
- We recommend that you use a static page for health checks. If the static page you use is not the default health check page of the backend ECS instance, you must set this page as the health check page in health check configurations. We recommend that you use a simple HTML page for health checks and use the page only for checking health check responses. We do not recommend that you use dynamic scripting languages such as php.
- Check whether the backend ECS instance has high loads, which can slow the response speed of the ECS instance.

Besides, because the layer-7 SLB service communicates with the backend ECS instance through an internal network, the ECS instance must listen to the internal network or all-network ports. You can check the ECS instance with the following methods:

1. Check whether the listening function is normal.

   Assume that the frontend port of SLB and backend port of the ECS instance are both 80. The ECS internal IP address is 10.11.192.1. Run the following command on the server. If you can see the monitoring information of 10.1.1.192.1: 80, or the monitoring information of 0.0.0.0: 80, the listening function of the ports is normal.

   - Windows server: `netstat -ano | findstr :80`
   - Linux server: `netstat -anp | grep :80`

2. Check whether the internal network firewall of the server allows port 80. You can disable the firewall temporarily to do a test. Enter the following command to disable the firewall.
   - Windows: `firewall.cpl`
   - Linux: `/etc/init.d/iptables stop`
3. Check whether the backend port is normal.
   - For layer-4 SLB service, you can perform a telnet test. If you receive responses, the backend port

is normal. Example: Use `telnet 10.11.192.1 80` to test.

- For layer-7 SLB service, you can determine whether the port is normal by checking the HTTP status code received. The HTTP status code must be a status code that indicates a normal condition, such as 200. The test methods are as follows:
  - Windows: Access the internal IP address of the ECS instance. In this example, access `http://10.11.192.1`.
  - Linux: Run the `curl -I` command and check whether the status is HTTP/1.1 200 OK. In this example, run `curl -I 10.11.192.1`.

5.How do I troubleshoot health check exceptions of a layer-4 (TCP/UDP) listener?

This topic describes how to troubleshoot a health check exception of a layer-4 (TCP/UDP) listener. The health check function is used to determine whether your backend servers are healthy. When a health check exception occurs, it generally means that your backend server is unhealthy. The exception may also be caused by incorrect health check configurations.

Procedure

1. Make sure that the backend server does not block the CIDR block 100.64.0.0/10 through iptables or other third-party firewalls or security software.

   Server Load Balancer (SLB) communicates with backend servers by using IP addresses in the reserved CIDR block 100.64.0.0/10. If the CIDR block is blocked, health check exceptions occur and SLB cannot work normally.

2. Run the **telnet** command to test the backend server.
   i. Log on to the SLB console and check the health check configurations.

   By default, the port of the backend server is used as the **Health Check Port**. You can also set the port manually. In this example, the port of the backend server, namely port 80, is used.

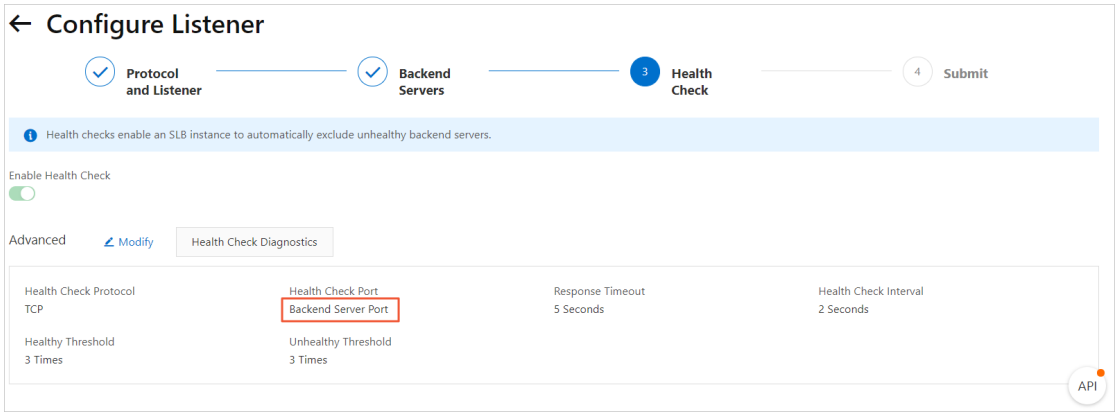

ii. Run the following command to connect the health check port. The health check port configured on the SLB instance must be the same as the listening port on the backend server.

```
telnet 172.17.58.131 80
```

In this example, *172.17.58.131* is the internal IP address of the backend server, and *80* is the health check port. By default, the port of the backend server is used as the health check port. You can configure the health check port according to your actual situation.

- In normal conditions, `Connected to xxx.xxx.xxx.xxx` is returned. This indicates that the port on the backend server is working (listening) normally and the health check succeeds, as shown in the following figure.

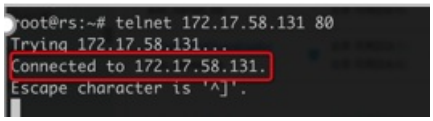

- Exception example: Assume you do not change the listener configurations of the SLB instance but stop the listening process of port 80 on the backend server. Then, if you run the **telnet** command, the system prompts that the host cannot be connected. This means that a health check exception occurs if the listening process of port 80 stops, as shown in the following figure.

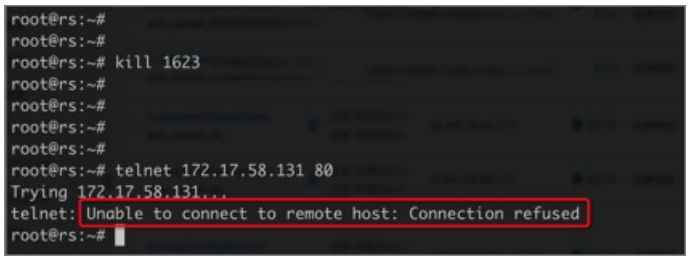

3. (Optional)Layer-4 listeners support HTTP health checks. If you use HTTP health checks, see How do I troubleshoot a health check exception of a layer-7 (HTTP/HTTPS) listener? for troubleshooting.

6.How do I troubleshoot a health check exception of a layer-7 (HTTP/HTTPS) listener?

This topic describes how to troubleshoot a health check exception of a layer-7 (HTTP/HTTPS) listener. The health check function is used to determine whether your backend servers are healthy. If a health check exception occurs, it generally means that your backend server is unhealthy. The exception may also be caused by incorrect health check configurations.

Procedure

1. Make sure that the backend server does not block the CIDR block 100.64.0.0/10 through iptables or other third-party firewalls or security software.

   Server Load Balancer (SLB) communicates with backend servers by using IP addresses in the reserved CIDR block 100.64.0.0/10. If the CIDR block is blocked, health check exceptions occur and SLB cannot work normally.

2. Access the HTTP service on the backend server from the backend server to check whether the HTTP service works normally.

   i. Log on to the SLB console and check the health check configurations on the listener details page.

   In this example, an HTTP listener is used and the internal IP address of the backend server with the health check exception is 10.0.0.2. Other health check configurations are as follows:

   - **Health Check Port**: 80
   - **Health Check Domain Name (Optional)**: `www.slb-test.com`
   - **Health Check Path**: `/test.html`

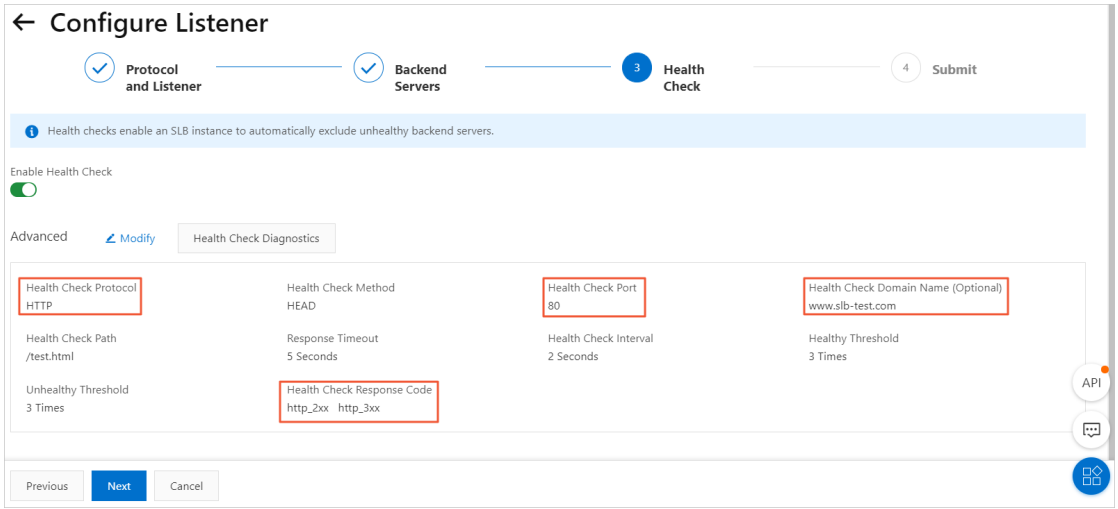

ii. For a Linux server, run the **nc** or **curl** command to test the HTTP service on the backend server. Make sure that the configurations of health check path, health check port, and health check domain name are the same for the HTTP service and the backend server. Otherwise, a health check exception occurs.

In this example, the **nc** command is used. Configure the health check path, health check domain name, public IP address, and health check port according to your actual situation.

```
echo -e "HEAD /test.html HTTP/1.0\r\nHost: www.slb-test.com\r\n\r\n" | nc -t 172.17
.58.131 80
```

- In normal conditions, `200` or `2xx/3xx` status codes are returned, as shown in the following figure.

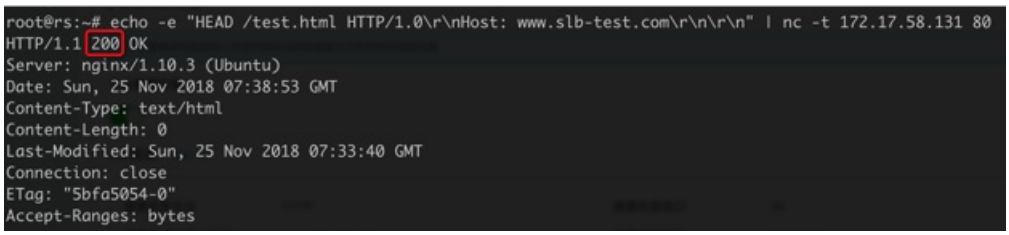

- Assume you do not change the listener configurations of the SLB instance but delete the /test.html page on the backend server. Then, when you run the **nc** command, the error code 404, instead of 2xx or 3xx, is returned, indicating a health check exception, as shown in the following figure.

7.How do I troubleshoot HTTP 5xx errors?

After a Server Load Balancer (SLB) instance is configured, errors such as 500 Internal Server Error, 502 Bad Gateway, and 504 Gateway Timeout may occur. These errors can be caused by the blockage from Internet service providers (ISPs), blockage from Alibaba Cloud Security for abnormal client activities, configuration errors of the SLB instance, health check failures, or failures in accessing web applications on the backend ECS instances.

This topic lists the causes, solutions, and troubleshooting methods of these problems.

Possible causes and solutions

1. The origin domain name does not have an ICP license or ICP filing, or no Layer 7 routing methods are configured for the domain name in the Alibaba Cloud anti-DDoS network or security network.

   Solution: Obtain an ICP filing or ICP license for the domain name? If the SLB instance is deployed in the Alibaba Cloud anti-DDoS network or security network, configure corresponding domain name-based routing methods.

2. The source IP address of the client is blocked by Alibaba Cloud Security.

   Check whether the same problem occurs to clients of other ISPs. If not, the problem is caused by blockage from the ISP.

   Solution: Submit a ticket to Alibaba Cloud technical support personnel, who then capture packets to determine whether the blockage occurs. If the blockage occurs, contact the ISP to solve the problem.

3. The source IP address is blocked by the security software of the backend ECS instance.

   100.64.0.0/10 is an CIDR block reserved by Alibaba Cloud for SLB servers for health checks and request forwarding. No security risks exist for 100.64.0.0/10. If security software or a firewall inside the system is applied, add this CIDR block to the whitelist of the software or firewall to avoid 500 or 502 errors.

Solution: Add 100.64.0.0/10 to the whitelist of antivirus or firewall software or uninstall the software to check whether the problem is caused by the blockage from the software.

4. Parameters of the Linux kernel of the backend ECS instance are incorrectly configured.

   If the backend ECS instance uses the Linux system, you can disable the rp_filter parameters in the system kernel when you change the Layer 7 listener to a Layer 4 listener.

   Solution: Set the values of the following parameters in the */etc/sysctl.conf* system configuration file to 0, and then run `sysctl -p`.

   ```
   net.ipv4.conf.default.rp_filter = 0
   net.ipv4.conf.all.rp_filter = 0
   net.ipv4.conf.eth0.rp_filter = 0
   ```

5. The backend ECS instance runs into a performance bottleneck.

   High CPU utilization or public bandwidth exhaustion may cause access exceptions.

   Solution: Check the performance of the backend ECS instance and solve performance bottlenecks. If the overall system capacity is insufficient, you can increase the number of backend ECS instances.

6. SLB reports 502 errors due to health check failures.

   For information about how to troubleshoot health check failures, see What can I do if my ECS instance is declared unhealthy after I enable health checks for Server Load Balancer?

   Also, 502 errors occur if the health check function of SLB is disabled and the web service in the backend server cannot process HTTP requests.

7. Health checks succeed but 502 errors are reported for web applications.

   If the 502 Bad Gateway error message indicates that SLB can forward requests from clients to backend servers, but the web applications on the backend servers cannot process the requests, you must check the configurations and running status of the web applications on the backend servers. For example, the time used by a web application to process an HTTP request exceeds the timeout value of SLB.

   For Layer 7 listeners, if the time used by the backend server to process PHP requests exceeds the proxy_read_timeout value of 60 seconds, SLB reports 504 Gateway Timeout. For Layer 4 listeners, the timeout value is 900 seconds.

   Solution: Make sure that the web service and related services run normally. Check whether PHP requests are processed properly, and optimize the processing of PHP requests by the backend server. NGINX+PHP-FPM is used in the following example:

   i. The number of PHP requests that are processed has reached the limit.

      When the total number of PHP requests that are processed in the server has reached the limit set by max_children in PHP-FPM, 502 or 504 erros may occur if more PHP requests are sent to the server:

      - If existing and new PHP requests are both processed in a timely manner, no error occurs.
      - New PHP requests wait to be processed when existing PHP requests are being processed. If the time that a new PHP request waits exceeds the value of fastcgi_read_timeout of NGINX, a 504 Gateway Timeout error occurs.
      - New PHP requests wait to be processed when existing PHP requests are being processed. If the time that a new PHP request waits exceeds the value of request_terminate_timeout of NGINX, a 502 Bad Gateway error occurs.

ii. If the PHP script execution time exceeds the limit, or the time used by PHP-FPM to process PHP scripts exceeds the value of request_terminate_timeout in NGINX, a 502 error occurs and the following error entry is shown in NGINX logs:

```
[error] 1760#0: *251777 recv() failed (104: Connection reset by peer) while reading response header from upstream, client: xxx.xxx.xxx.xxx, server: localhost, request: "GET /timeoutmore.php HTTP/1.1", upstream: "fastcgi://127.0.0.1:9000"
```

iii. The health check is performed on static pages. Errors occur when exceptions are detected in the process that handles dynamic requests. For example, PHP-FPM is not running.

8. The HTTP header is too long.

   An HTTP header that is too long may make SLB unable to process relevant data, which results in 502 errors.

   Solution: Decrease the amount of data transmitted by the header or switch to the TCP listener.

9. The service access logic is inappropriate.

   Make sure that no backend ECS instance in SLB accesses the public IP addresses of SLB instances. If a backend ECS instances accesses its own port through the public IP address of the SLB instance to which the ECS instance is added as a backend server, the access request is sent to the backend ECS instance based on the routing methods configured for the SLB instance. This leads to an infinite loop, which results in 500 or 502 errors.

   Solution: Make sure that no backend ECS instance in SLB accesses the public IP addresses of SLB instances.

Troubleshooting

- Check the screenshot of the 500, 502, or 504 error to determine the cause of the error. The error may be caused by SLB, Anti-DDoS or security network, or backend ECS instance configurations.
- If Anti-DDoS or security network is used, make sure that the Layer 7 routing methods are correctly configured.
- Check whether the problem occurs to all clients. If not, check whether the client that indicates an error has been blocked by Alibaba Cloud Security. Also, check whether the domain name or IP address of SLB is clocked by the ISP.
- Check the status of SLB instances and whether backend ECS instances fail the health check. If one or more backend ECS instances fail the health check, troubleshoot the failure.
- Associate the endpoint of SLB with the IP address of the backend server by using the hosts file on the client. If a 5xx error occurs at intervals, the error may be caused by incorrect configurations of a backend ECS instance.
- Change Layer 7 SLB to Layer 4 SLB to see whether the problem occurs again.
- Check the performance of backend ECS servers and whether performance bottlenecks of the CPU, memory, disk, or bandwidth exist.
- If the error is caused by the backend server, check the web server logs of the backend server. Check whether the web service is running normally and whether the web access logic is correct.
- Check whether the TCP kernel parameters of the Linux system on the backend ECS instance are correctly configured.

Submit a ticket

Perform the preceding troubleshooting procedures step by step and record the test results in detail. Provide the test results when you submit a ticket so that Alibaba Cloud technical support personnel can help you solve the problem promptly.

If the problem persists, consult Alibaba Cloud technical support personnel.

8.Configure cookie in the backend server

Server Load Balancer provides session persistence function. With session persistence enabled, Server Load Balancer can distribute requests from the same client to the same backend server during the session period.

For layer-4 listeners, session persistence is based on the IP address. The listener of Server Load Balancer forwards requests from the same IP address to the same backend server. For layer-7 listeners, session persistence is based on cookies.

If you choose to rewrite the cookie, you must configure the cookie on the backend server. Suppose there are two domain names under your Server Load Balancer service: vip.a.com and img.a.com. If you want to configure session persistence for vip.a.com, you can set the cookie name to name, and set a cookie of which the key is name for vip.a.com on the backend server.

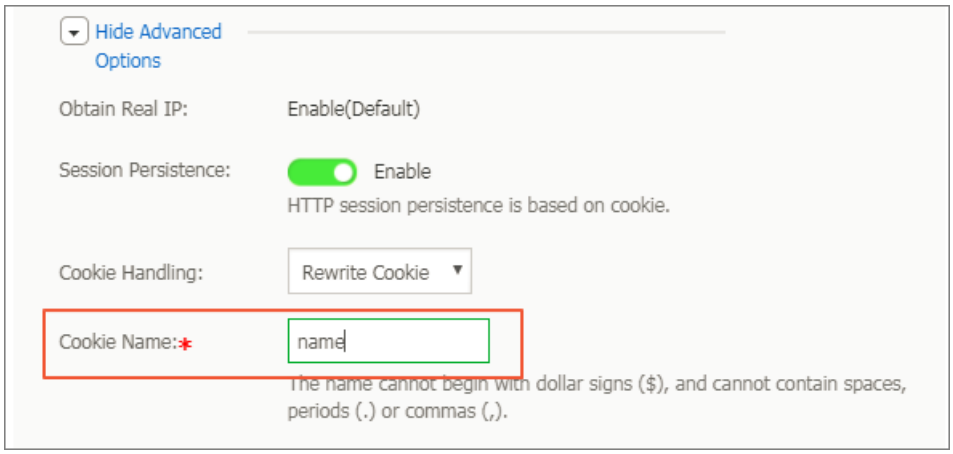

Follow the instructions in this section to set cookies on a backend server.

Apache

1. Open the *httpd.conf* file and make sure that the following line is not commented.

   ```
   LoadModule usertrack_module modules/mod_usertrack.so
   ```

2. Add the following configurations in the VirtualHost file.

   ```
   CookieName name
   CookieExpires "1 days"
   CookieStyle Cookie
   CookieTracking on
   ```

Nginx

Configure the cookie as follows.

```
server {
    listen 8080;
    server_name wqwq.example.com;
    location / {
      add_header Set-Cookie name=xxxx;
        root html;
        index index.html index.htm;
    }
}
```

Lighttpd

Configure the cookie as follows.

```
    server.modules  = ( "mod_setenv" )
    $HTTP["host"] == "test.example.com" {
          server.document-root = "/var/www/html/"
          setenv.add-response-header = ( "Set-Cookie" => "name=XXXXXX"       }
    }
```

9.FAQ about session persistence

What is session persistence used for?

Session persistence is a method to forward requests from the same client to the same backend server.

How do I enable session persistence?

You can enable session persistence when you configure a

Classic Load Balancer (CLB)

listener. You can configure different session persistence policies for different listeners. A session can be persisted for at most 86400 seconds, which is equivalent to 24 hours.

What types of session persistence does

CLB

support?

- CLB

  persists TCP and UDP sessions based on source IP addresses at Layer 4. A session can be persisted for at most 3600 seconds at Layer 4.

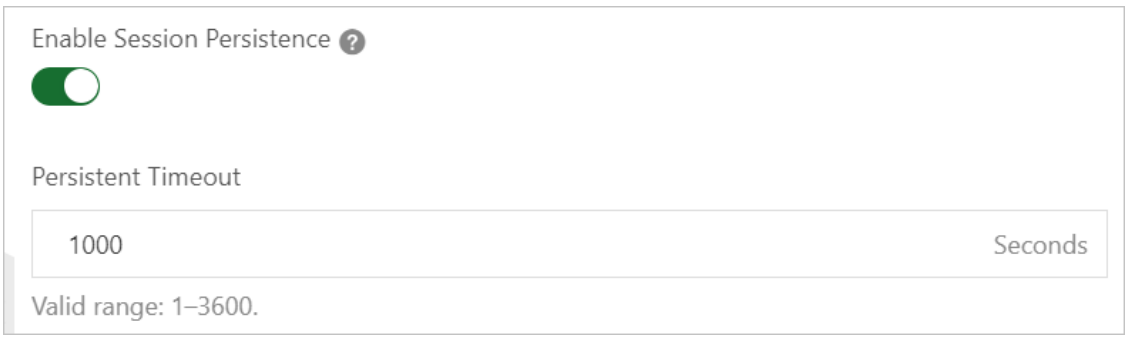

- CLB

  persists HTTP and HTTPS sessions based on cookies at Layer 7. A cookie-based session can be persisted for at most 86400 seconds, which is equivalent to 24 hours.

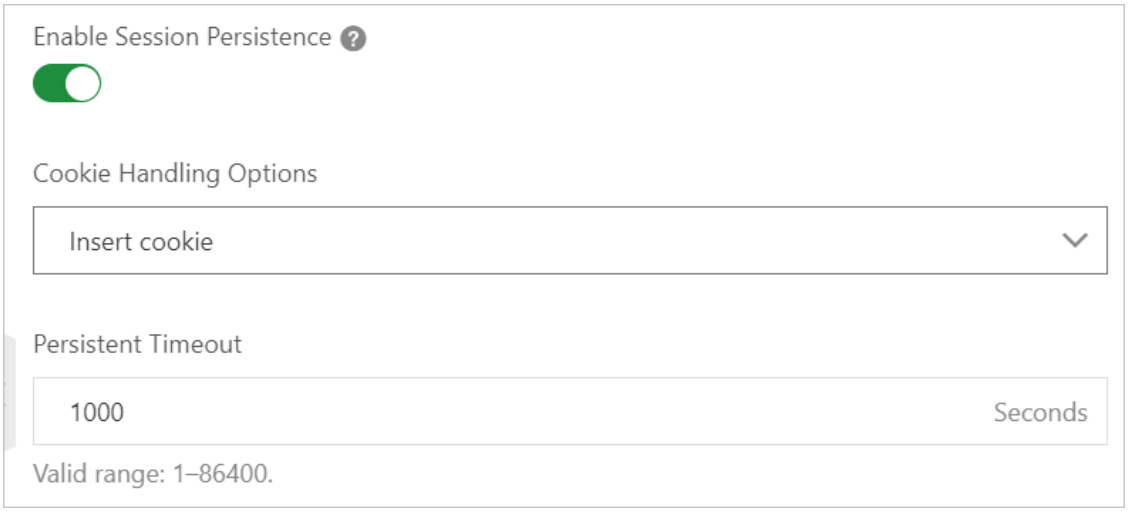

What are the methods that can be used to handle cookies?

HTTP and HTTPS listeners support **cookie inserts** and **cookie rewrites**.

- **Insert a cookie**: If you select this option, you need only to specify the timeout period of the cookie.

  CLB

  inserts a cookie (SERVERID) into the first HTTP or HTTPS response packet that is sent to a client. The next request from the client contains this cookie, and

  CLB

  forwards this request to the recorded Elastic Compute Service (ECS) instance.
- **Rewrite a cookie**: If you select this option, you can specify the cookie to be carried in an HTTP or HTTPS response. You must configure the timeout period and the lifetime of a cookie on an ECS instance. When

  CLB

  detects a user-defined cookie, CLB overwrites the original cookie with the user-defined cookie. The next request from the client carries the user-defined cookie, and

  CLB

  forwards this request to the recorded backend server. For more information about how to configure cookies on a server, see Configure session persistence.

Can CLB persist sessions based on domain names?

Yes, CLB can persist sessions based on domain names.

CLB

can rewrite cookies to persist sessions based on domain names.

How long is the timeout period of a cookie?

- You can specify the timeout period of a cookie (SERVERID) to be inserted by CLB from 1 to 86400 seconds in the console.

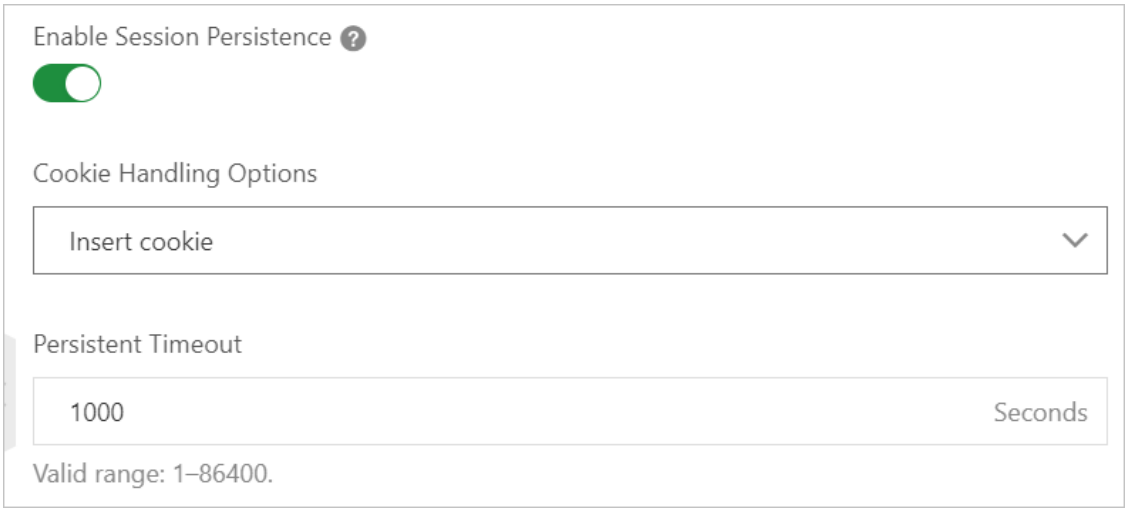

- For a user-defined cookie with which you want CLB to overwrite, you must configure the timeout period on the ECS instance.

How do I view a cookie?

Open the browser and press F12 to check whether SERVERID or a user-defined cookie is inserted in the response. You can also run the `curl www.example.com -c /tmp/cookie123` command to save a cookie and then run the `curl www.example.com -b /tmp/cookie123` command to view the cookie.

Why does session persistence fail?

- Check whether session persistence is enabled for the listener.
- HTTP and HTTPS listeners cannot persist sessions by inserting cookies into responses that carry 4xx status codes.

  Solution: Use TCP listeners instead of HTTP or HTTPS listeners. TCP listeners persist sessions based on client IP addresses. Backend servers can also insert or even validate cookies to ensure that sessions are persisted.
- HTTP 302 redirects change the SERVERID string for persisting a session.

  When CLB inserts a cookie into a response that carries the HTTP status code 302, the SERVERID string is changed. As a result, the session cannot be persisted.

  To verify the cause, check the requests and responses by using your browser or packet capture software. Then, check whether a 302 status code is included in the packets and whether the SERVERID string in the cookie is changed.

  Solution: Use TCP listeners instead of HTTP or HTTPS listeners. TCP listeners persist sessions based on client IP addresses. Backend servers can also insert or even validate cookies to ensure that sessions are persisted.
- The timeout period is set to a small value. You can set the timeout period to a greater value.

How do I verify session persistence by using the Linux curl command?

1. Create a test page.

Create a test page on each backend server. You can view the private IP address of the backend server on the test page. The following figure shows an example of a test page. The private IP address indicates the backend server to which requests are distributed. The private IP address is used to check whether CLB can persist sessions.

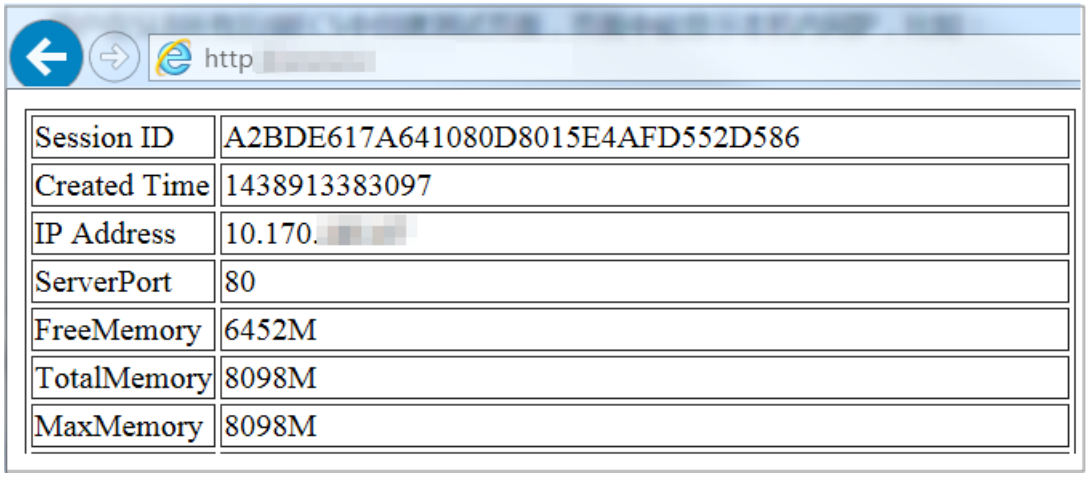

2. Run the curl command in Linux.

   In this example, the IP address of the CLB instance that runs Linux is 10.170.XX.XX and the URL of the created page is `http://10.170.XX.XX/check.jsp`.

   i. Log on to a server that runs Linux.

   ii. Run the following command to query the cookie inserted by the backend server:

   ```
   curl -c test.cookie http://10.170.XX.XX/check.jsp
   ```

   > **Note** By default, CLB persists sessions by inserting cookies. However, curl does not send or save cookies. Therefore, you must save a cookie before you perform the test. Otherwise, the curl test result may show that session persistence is invalid.

   iii. After you save the cookie, run the following command:

   ```
   for ((a=1;a<=30;a++));
       do curl  -b test.cookie http://10.170.XX.XX/check.jsp  | grep '10.170.XX.XX';
       sleep 1;
   done
   ```

   > **Note** a<=30 indicates the number of tests to be performed. You can set this value based on your business requirements. Set the IP address in `grep '10.170.XX.XX'` to the private IP address of your ECS instance.

   iv. Check the IP addresses returned in the preceding tests. If the same IP address is returned, it indicates that CLB can persist sessions.