



Configuring Health Monitoring

This chapter describes how to configure the health monitoring on the CSM-S and contains these sections:

- [Configuring Probes for Health Monitoring, page 11-1](#)
- [Understanding and Configuring Inband Health Monitoring, page 11-8](#)
- [Understanding and Configuring HTTP Return Code Checking, page 11-9](#)

Configuring Probes for Health Monitoring

To efficiently balance connections, the CSM must continuously monitor the health of all real servers in its configuration. Configuring health probes to the real servers allows you to determine if the real servers are operating correctly. A real server's health is categorized as follows:

- Active—The real server responds appropriately.
- Suspect—The real server is unreachable or returns an invalid response. The probes are retried.
- Failed—The real server fails to reply after a specified number of consecutive retries. You are notified and the CSM-S adjusts incoming connections accordingly. Probes continue to a failed server until the server becomes active again.

The CSM-S supports probes used to monitor real servers. Configuring a probe involves the following:

- Entering the probe submode
- Naming the probe
- Specifying the probe type

The CSM-S supports a variety of probe types that monitor real servers, including FTP, DNS, or HTTP.



Note

By default, no probes are configured on the CSM-S.

When configuring the CSM-S for health probe monitoring, you can use a multiple-tiered approach that includes the following actions:

- Active probes—These probes run periodically. ICMP, TCP, HTTP, and other predefined health probes fall into this category. Scripted health probes are included here as well. Active probes do not impact the session setup or teardown system.

- Passive monitoring (in-band health monitoring)—Monitors sessions for catastrophic errors that can remove a server from services. Catastrophic errors may be reset (RST) when there is no response from a server. These health checks operate at a full-session rate, and recognize failing servers quickly.
- Passive HTTP error code checking (in-band response parsing)—The CSM-S parses HTTP return codes and watches for codes such as “service unavailable” so that it can take a server out of service. Passive HTTP error code checking has a small impact on session performance.

To set up a probe, you must configure it by naming the probe and specifying the probe type while in probe submenu.

After configuring a probe, you must associate it with a server farm for the probe to take effect. All servers in the server farm receive probes of the probe types that are associated with that server farm. You can associate one or more probe types with a server farm.

If you assign a port number when configuring either the real server or the virtual server, you do not need to specify a port number when you configure a probe. The probe inherits the port number from the real or virtual server configuration.

You can override the real server’s and virtual server’s port information by explicitly specifying a port to probe in the health probe configuration using the optional health probe port feature. This feature allows you to set a port for use by the health probes when no port is specified either in the real server or virtual server.

After you configure a probe, associate single or multiple probes with a server farm. All servers in the server farm receive probes of the probe types that are associated with that pool.



Note

If you associate a probe of a particular type with a server farm containing real servers that are not running the corresponding service, the real servers send error messages when they receive a probe of that type. This action causes the CSM-S to place the real server in a failed state and disable the real server from the server farm.

To specify a probe type and name, perform this task:

	Command	Purpose
Step 1	<pre>Router(config-module-csm)# probe probe-name {http icmp telnet tcp ftp smtp dns kal-ap-upd}</pre>	<p>Specifies a probe type and a name^{1 2}:</p> <ul style="list-style-type: none"> • <i>probe-name</i> is the name of the probe being configured; it has a character string of up to 15 characters. • http creates an HTTP probe with a default configuration. • icmp creates an ICMP probe with a default configuration. • telnet creates a Telnet probe with a default configuration. • tcp creates a TCP probe with a default configuration. • ftp creates an FTP probe with a default configuration. • smtp creates an SMTP probe with a default configuration. • dns creates a DNS probe with a default configuration. • kal-ap-upd creates a GSLB target probe.
Step 2	<pre>Router(config-slb-probe-tcp)# port port-number: 1-MAXUSHORT</pre>	<p>Configures an optional port for a probe³.</p>

	Command	Purpose
Step 3	Router# show module csm slot probe	Displays all probes and their configuration.
Step 4	Router# show module csm slot tech-support probe	Displays probe statistics.

1. The **no** form of this command removes the probe type from the configuration.
2. Inband health monitoring provides a more scalable solution if you are receiving performance alerts.
3. The **port** command does not exist for the ICMP or PING health probe.

**Note**

When you specify a probe name and type, it is initially configured with the default values. Enter the probe configuration commands to change the default configuration.

This example shows how to configure a probe:

```
Router(config-module-csm)# probe probe1 tcp
Router(config-slb-probe-tcp)# interval 120
Router(config-slb-probe-tcp)# retries 3
Router(config-slb-probe-tcp)# failed 300
Router(config-slb-probe-tcp)# open 10
Router(config-slb-probe-tcp)# serverfarm sf4
Router(config-slb-sfarm)# real 10.1.0.105
Router(config-slb-real)# inservice
Router(config-slb-real)# probe probe1
Router(config-slb-sfarm)# vserver vs4
Router(config-slb-vserver)# virtual 10.1.0.84 tcp 80
Router(config-slb-vserver)# serverfarm sf4
Router(config-slb-vserver)# inservice
Router(config-slb-vserver)# end
```

**Note**

There are two different timeout values: open and receive. The open timeout specifies how many seconds to wait for the connection to open (that is, how many seconds to wait for SYN ACK after sending SYN). The receive timeout specifies how many seconds to wait for data to be received (that is, how many seconds to wait for an HTTP reply after sending a GET/HHEAD request). Because TCP probes close as soon as they open without sending any data, the receive timeout is not used.

Probe Configuration Commands

These commands are common to all probe types:

Command	Purpose
Router (config-slb-probe) # interval <i>seconds</i>	Sets the interval between probes in seconds (from the end of the previous probe to the beginning of the next probe) when the server is healthy ¹ ² . Range = 2–65535 seconds Default = 120 seconds
Router (config-slb-probe) # retries <i>retry-count</i>	Sets the number of failed probes that are allowed before marking the server as failed ¹ . Range = 0–65535 Default = 3
Router (config-slb-probe) # failed <i>failed-interval</i>	Sets the time between health checks when the server has been marked as failed. The time is in seconds ¹ . Range = 2–65535 Default = 300 seconds
Router (config-slb-probe) # recover <i>recover-value</i>	Sets the number of consecutive responses that are received before marking a failed server as healthy ¹ . Range = 1–65535 Default = 1
Router (config-slb-probe) # open <i>open-timeout</i>	Sets the maximum time to wait for a TCP connection. This command is not used for any non-TCP health checks (ICMP or DNS ¹). Range = 1–65535 Default = 10 seconds
Router (config-slb-probe) # description <i>description</i>	(Optional) Specifies a description for the probe. Limit the <i>description</i> to 80 characters.

1. The **no** form of this command restores the defaults.
2. Inband health monitoring provides a more scalable solution if you are receiving performance alerts.

Configuring an HTTP Probe

An HTTP probe establishes an HTTP connection to a real server and then sends an HTTP request and verifies the response. The **probe** *probe-name* **http** command places the user in HTTP probe configuration submode.

To configure an HTTP probe, perform this task:

	Command	Purpose
Step 1	Router (config-module-csm) # probe <i>probe-name</i> http	Configures an HTTP probe and enters the HTTP probe submode ¹ .
Step 2	Router (config-slb-probe-http) # credentials <i>username</i> [<i>password</i>]	Configures basic authentication values for the HTTP SLB probe ¹ .

	Command	Purpose
Step 3	Router(config-slb-probe-http)# expect status <i>min-number</i> [<i>max-number</i>]	<p>Configures a status code to expect from the HTTP probe. You can configure multiple status ranges by entering one expect status command at a time¹.</p> <p><i>min-number</i>—If you do not specify a <i>max-number</i>, this number is taken as a single status code. If you specify a maximum number, this number is taken as the minimum status code of a range.</p> <p><i>max-number</i>—The maximum status code in a range. The default range is 0–999. (Any response from the server is considered valid.)</p> <p>Note If no maximum is specified, this command takes a single number (<i>min-number</i>). If you specify both a minimum number and a maximum number, it takes the range of numbers.</p>
Step 4	Router(config-slb-probe-http)# header <i>field-name</i> [<i>field-value</i>]	Configures a header field for the HTTP probe. Multiple header fields may be specified ¹ .
Step 5	Router(config-slb-probe-http)# request [<i>method</i> [get head]] [<i>url path</i>]	<p>Configures the request method used by an HTTP probe¹:</p> <ul style="list-style-type: none"> • get—The HTTP get request method directs the server to get this page. • head—The HTTP head request method directs the server to get only the header for this page. • url—A character string of up to 1275 characters specifies the URL path; the default path is “/”. <p>Note The CSM-S supports only the get and head request methods; it does not support the post and other methods. The default method is get.</p>

1. The **no** form of this command restores the defaults.

Configuring an ICMP Probe

An ICMP probe sends an ICMP echo (for example, ping) to the real server. The **probe icmp** command enters the ICMP probe configuration mode. All the common **probe** commands are supported except the **open** command, which is ignored.

To configure an ICMP probe, perform this task:

	Command	Purpose
Step 1	Router(config-module-csm)# probe <i>probe-name</i> icmp	Configures an ICMP probe and enters the ICMP probe submode ¹ .
Step 2	Router(config-slb-probe-icmp)# interval <i>seconds</i>	Configures the intervals to wait between probes of a failed server and between probes.

	Command	Purpose
Step 3	Router(config-slb-probe-icmp)# receive <i>seconds</i>	Specifies the time to make a TCP connection to receive a reply from the server.
Step 4	Router(config-slb-probe-icmp)# retries <i>count</i>	Limits the number of retries before considering the server as failed.

1. The **no** form of this command restores the defaults.

Configuring a UDP Probe

The UDP probe requires ICMP because otherwise the UDP probe will be unable to detect when a server has gone down or has been disconnected. You must associate UDP to the supervisor engine and then configure ICMP.

Because the UDP probe is a raw UDP probe, the CSM-S is using a single byte in the payload for probe responses. The CSM-S does not expect any meaningful response from the UDP application. The CSM-S uses the ICMP Unreachable message to determine if the UDP application is not reachable. If there is no ICMP unreachable reply in the receive timeout, the CSM-S assumes that the probe is operating correctly. If the IP interface of the real server is down or disconnected, the UDP probe by itself would not know that the UDP application is not reachable. You must configure the ICMP probe in addition to the UDP probe for any given server.

The CSM-S uses the DNS probe as the high-level UDP application. You can use a TCL script to configure this probe. See [Chapter 12, “Using TCL Scripts with the CSM-S.”](#)

To configure an ICMP probe, perform this task:

	Command	Purpose
Step 1	Router(config-module-csm)# probe <i>probe-name</i> udp	Configures a UDP probe and enters the UDP probe submenu ¹ .
Step 2	Router(config-slb-probe-icmp)# interval <i>seconds</i>	Configures the intervals to wait between probes of a failed server and between probes.
Step 3	Router(config-slb-probe-icmp)# receive <i>seconds</i>	Specifies the time to make a TCP connection to receive a reply from the server.
Step 4	Router(config-slb-probe-icmp)# retries <i>count</i>	Limits the number of retries before considering the server as failed.

1. The **no** form of this command restores the defaults.

Configuring a TCP Probe

A TCP probe establishes and removes connections. The **probe tcp** command enters the TCP probe configuration mode. All the common **probe** commands are supported.

To configure a TCP probe, perform this task:

	Command	Purpose
Step 1	Router(config-module-csm) # probe <i>probe-name</i> tcp	Configures a TCP probe and enters the TCP probe submode ¹ .
Step 2	Router(config-slb-probe-icmp) # interval <i>seconds</i>	Configures the intervals to wait between probes of a failed server and between probes.
Step 3	Router(config-slb-probe-icmp) # retries <i>count</i>	Limits the number of retries before considering the server as failed.

1. The **no** form of this command restores the defaults.

Configuring FTP, SMTP, and Telnet Probes

An FTP, SMTP, or Telnet probe establishes a connection to the real server and verifies that a greeting from the application was received. The **probe** (**ftp**, **smtp**, or **telnet**) command enters the corresponding probe configuration mode. All the **probe** common options are supported. Multiple status ranges are supported, one command at a time.

To configure a status code to expect from the FTP, SMTP, or Telnet probe, perform this task:

	Command	Purpose
Step 1	Router(config-module-csm) # probe <i>probe-name</i> [ftp smtp telnet]	Configures an FTP, SMTP, or Telnet probe and enters the FTP, SMTP, or Telnet probe submode ¹ .
Step 2	Router(config-slb-probe-icmp) # interval <i>seconds</i>	Configures the intervals to wait between probes of a failed server and between probes.
Step 3	Router(config-slb-probe-icmp) # receive <i>seconds</i>	Specifies the time to make a TCP connection to receive a reply from the server.
Step 4	Router(config-slb-probe-icmp) # retries <i>count</i>	Limits the number of retries before considering the server as failed.

1. The **no** form of this command restores the defaults.

Specifying the DNS Resolve Request

A DNS probe sends a domain name resolve request to the real server and verifies the returned IP address. The **probe dns** command places the user in DNS probe configuration submode. All the probe common options are supported except **open**, which is ignored.

To specify the domain name resolve request, perform this task:

	Command	Purpose
Step 1	Router(config-module-csm)# probe <i>probe-name</i> dns	Configures a DNS probe and enters the tcp probe submode ¹ .
Step 2	Router(config-slb-probe-dns)# [failed interval retries receive] <i>number</i>	Configures the times to wait between probes to make a DNS connection, to receive a reply from the server, and to limit the number of retries before considering the real server as failed.

1. The **no** form of this command restores the defaults.

Configuring GSLB Probes

GSLB capabilities for the Cisco CSM were introduced in the CSM software release 3.1. With this capability, the CSM combines GSLB with route health injection.

When configuring Global Server Load Balancing (GSLB) type probes, the **port** submode command is not used to specify which destination UDP port to query. Use the CSM environment variable `GSLB_KALAP_UDP_PORT` instead. The default is port 5002.

To specify probe frequency and the number of retries for KAL-AP, ICMP, HTTP, and DNS probes when associated with a GSLB server farm environment, the following variables must be used instead of the probe configuration submode commands:

<code>GSLB_KALAP_PROBE_FREQ</code>	10
<code>GSLB_KALAP_PROBE_RETRIES</code>	3
<code>GSLB_ICMP_PROBE_FREQ</code>	10
<code>GSLB_ICMP_PROBE_RETRIES</code>	3
<code>GSLB_HTTP_PROBE_FREQ</code>	10
<code>GSLB_HTTP_PROBE_RETRIES</code>	2
<code>GSLB_DNS_PROBE_FREQ</code>	10
<code>GSLB_DNS_PROBE_RETRIES</code>	3

GSLB for the CSM is deployed such that the CSM becomes the DNS authoritative name server for the GSLB subdomain. Although the CSM does not use CAPP, it keeps track of local and remote devices by using health-checking probes. These probes are sent to the local real server within the data center and to the remote virtual server on the remote CSMs. The CSM should be configured with the same probes to monitor the availability of the local server and the remote virtual server at the active data center.

When DNS requests are sent to the CSM, the CSM assesses the load and availability of the local virtual servers and the remote virtual servers that support the requested domain. The CSM then responds with the IP address of the local or remote virtual servers, depending on the availability and load of these servers. If the information collected by the health probes confirms that the local virtual server is unavailable or overloaded, the CSM responds with the IP address of the remote virtual server.

Understanding and Configuring Inband Health Monitoring

These sections describe inband health monitoring:

- [Understanding Inband Health Monitoring, page 11-9](#)
- [Configuring Inband Health Monitoring, page 11-9](#)

Understanding Inband Health Monitoring

Inband health monitoring (passive monitoring) monitors normal traffic sessions for catastrophic errors (for example, a reset (RST) message from the server or no response from a server) that can remove a server from service.

**Note**

Inband health monitoring is not suitable for all types of services typically provided by server farms. In particular, failures of FTP and RTSP servers are not detected because these protocols require special handling that is performed outside of the processor that performs inbound health monitoring. Use other types of probes for these services.

The inband health monitoring feature is configured for each server farm to monitor the health of the servers. The parameters configured per server farm are then applied to each real server in that server farm. You can configure the number of abnormal end sessions that occur before the system considers the real server unreachable. You also can specify a time to wait before a real server is reintroduced into the server farm and a connection attempt is made.

This feature works with health probes. If health probes and inband health monitoring are both configured on a particular server, both sets of health checks are required to keep a real server in service within a server farm. If either health-checking feature finds a server out of service, the server will not be selected by the CSM-S for load balancing.

Configuring Inband Health Monitoring

To configure inband health monitoring, perform these steps:

- Step 1** Verify that you have configured server farms. (See the [“Configuring Server Farms”](#) section on page 5-1.)
- Step 2** Enter the `serverfarm` submode command to enable inband health monitoring for each server farm:

```
Router(config-module-csm)# serverfarm serverfarm-name  
Router(config-slb-sfarm)# health retries count failed seconds
```

**Note**

Retries are the number of abnormal end sessions that the CSM-S will tolerate before removing a real server from service. The failed time is the number of seconds that the CSM-S waits before reattempting a connection to a real server that was removed from service by inband health checking.

This example shows how to enable inband health monitoring for a server farm named `geo`:

```
Router(config-module-csm)# serverfarm geo  
Router(config-slb-sfarm)# health retries 43 failed 160
```

Understanding and Configuring HTTP Return Code Checking

These sections describe HTTP return code checking:

- [Understanding HTTP Return Code Checking, page 11-10](#)
- [Configuring HTTP Return Code Checking, page 11-10](#)

Understanding HTTP Return Code Checking

The return error code checking (return code parsing) feature is used to indicate when a server is not returning web pages correctly. This feature extends the capability of CSM-S to inspect packets, parse the HTML return codes, and act upon the return codes returned by the server.

After receiving an HTTP request from the CSM-S, the server responds with an HTTP return code. The CSM-S can use the HTTP return error codes to determine the availability of the server. The CSM-S can be configured to take a server out of use in response to receiving specific return codes.

A list of predefined codes (100 through 599) are in RFC 2616. For return code checking, some codes are more usable than others. For example, a return code of 404 is defined as a URL not found, which may be the result of the user entering the URL incorrectly. Error code 404 also might mean that the web server has a hardware problem, such as a defective disk drive preventing the server from finding the data requested. In this case, the web server is still alive, but the server cannot send the requested data because of the defective disk drive. Because of the inability of the server to return the data, you do not want future requests for data sent to this server. To determine the error codes you want to use for return code checking, refer to RFC 2616.

When HTTP return code checking is configured, the CSM-S monitors HTTP responses from all balanced HTTP connections and logs the occurrence of the return code for each real server. The CSM-S stores return code counts. When a threshold for a return code is reached, the CSM-S may send syslog messages or remove the server from service.

A default action, counting return codes, syslog messaging, or removing the real server from service or a set of these actions can be applied to a server farm. You also can bind a single virtual group to multiple server farms allowing you to reuse a single return code server farm policy on multiple server farms.



Note

When you configure HTTP return code checking on a virtual server, the performance of that virtual server is impacted. Once return code parsing is enabled, all HTTP server responses must be parsed for return codes.

Configuring HTTP Return Code Checking

When you configure return error code checking, you configure the attributes of a server farm and associate it with a return code map.

To configure the return code checking, perform these steps:

-
- Step 1** Verify that you have configured HTTP virtual servers. (See the [“Configuring Redirect Virtual Servers” section on page 6-6.](#))
- Step 2** Enter the map return code command to enable return code mapping and enter the return code map submode:
- ```
Router(config-module-csm)# map name retcode
```
- Step 3** Configure the return code parsing:
- ```
Router(config-slb-map-retcode)# match protocol http retcode min max action [count | log | remove] threshold [reset seconds]
```

You can set up as many matches as you want in the map.

Step 4 Assign a return code map to a server farm:

```
Router(config-slb-sfarm)# retcode-map name
```

This example shows how to enable return error code checking:

```
Router(config-module-csm)# map httpcodes retcode
Route(config-slb-map-retcode)# match protocol http retcode 401 401 action log 5 reset 120
Route(config-slb-map-retcode)# match protocol http retcode 402 415 action count
Route(config-slb-map-retcode)# match protocol http retcode 500 500 action remove 3 reset 0
Route(config-slb-map-retcode)# match protocol http retcode 503 503 action remove 3 reset 0
Route(config-slb-map-retcode)# exit
Router(config-module-csm)# serverfarm farm1
Router(config-slb-sfarm)# retcode-map httpcodes
Router(config-slb-sfarm)# exit
Router(config-module-csm)# end
```

