
Chapter 14

Route Health Injection

You can configure an HP 9304M, HP 9308M, and HP 6208M-SX routing switch to check the health of the HTTP application and “inject” a host route into the network to force a preferred route to an actively responding web host. The web host can be directly attached to the routing switch or can be attached through Layer 2 switches. The web host can be a web server or a Server Load Balancing (SLB) device configured with a virtual IP address (VIP) representing the HTTP application.

The **route health injection** feature enables a routing switch to advertise a host route to a globally-distributed web site. Gateway routers that receive the host route along with other routes to the same web site in other locations can choose the best route. Web clients attached to the gateway servers thus enjoy fast response time regardless of their location, because their gateway routers use the best path to the web site. By advertising the host route instead of a network route to the web site’s IP address, the routing switch ensures that gateway routers receive a route to the IP address only if that IP address is available. The routing switch uses a Layer-4 HTTP health check that you configure to determine whether the HTTP (web) service on the IP address is available. The health check and how to configure it are described later in this section.

NOTE: This feature supports health checks only for TCP port 80 (HTTP).

Normally, an IP address should exist on only one host on the public Internet. However, some third-party SLBs allow the same IP address to exist on multiple machines using virtual IP addresses (VIPs). A VIP is an IP address that you configure on a third-party SLB, then associate with “real” servers attached to the SLB. These real servers are the web hosts that contain the web site requested by clients. In a simple SLB configuration, a single SLB contains a VIP that maps to multiple real servers that have identical contents. The VIP is the IP address associated with the web site on DNS servers. In a globally-distributed SLB configuration, multiple SLBs in different networks throughout the Internet are configured with same VIP and are attached to sets of real servers that contain the web site's content.

Configuration Example

Suppose you configure an SLB in Los Angeles and another one in New York to serve VIP 209.157.22.249. For this example, also assume that you have a real server in Paris with the same IP address and the server is directly attached to a routing switch.

Suppose the DNS entry for this IP address maps the address to a site named `www.acmeweb.com`. When a web client in Los Angeles enters this domain in their web browser, the web browser goes to the client's local DNS to resolve the name into an IP address. When the DNS returns the address to the web browser, the browser then attempts to contact the HTTP port (usually TCP port 80) on the host with the IP address returned by the DNS.

Figure 14.1 shows an example of a globally-distributed SLB configuration in which the route health injection feature is used.

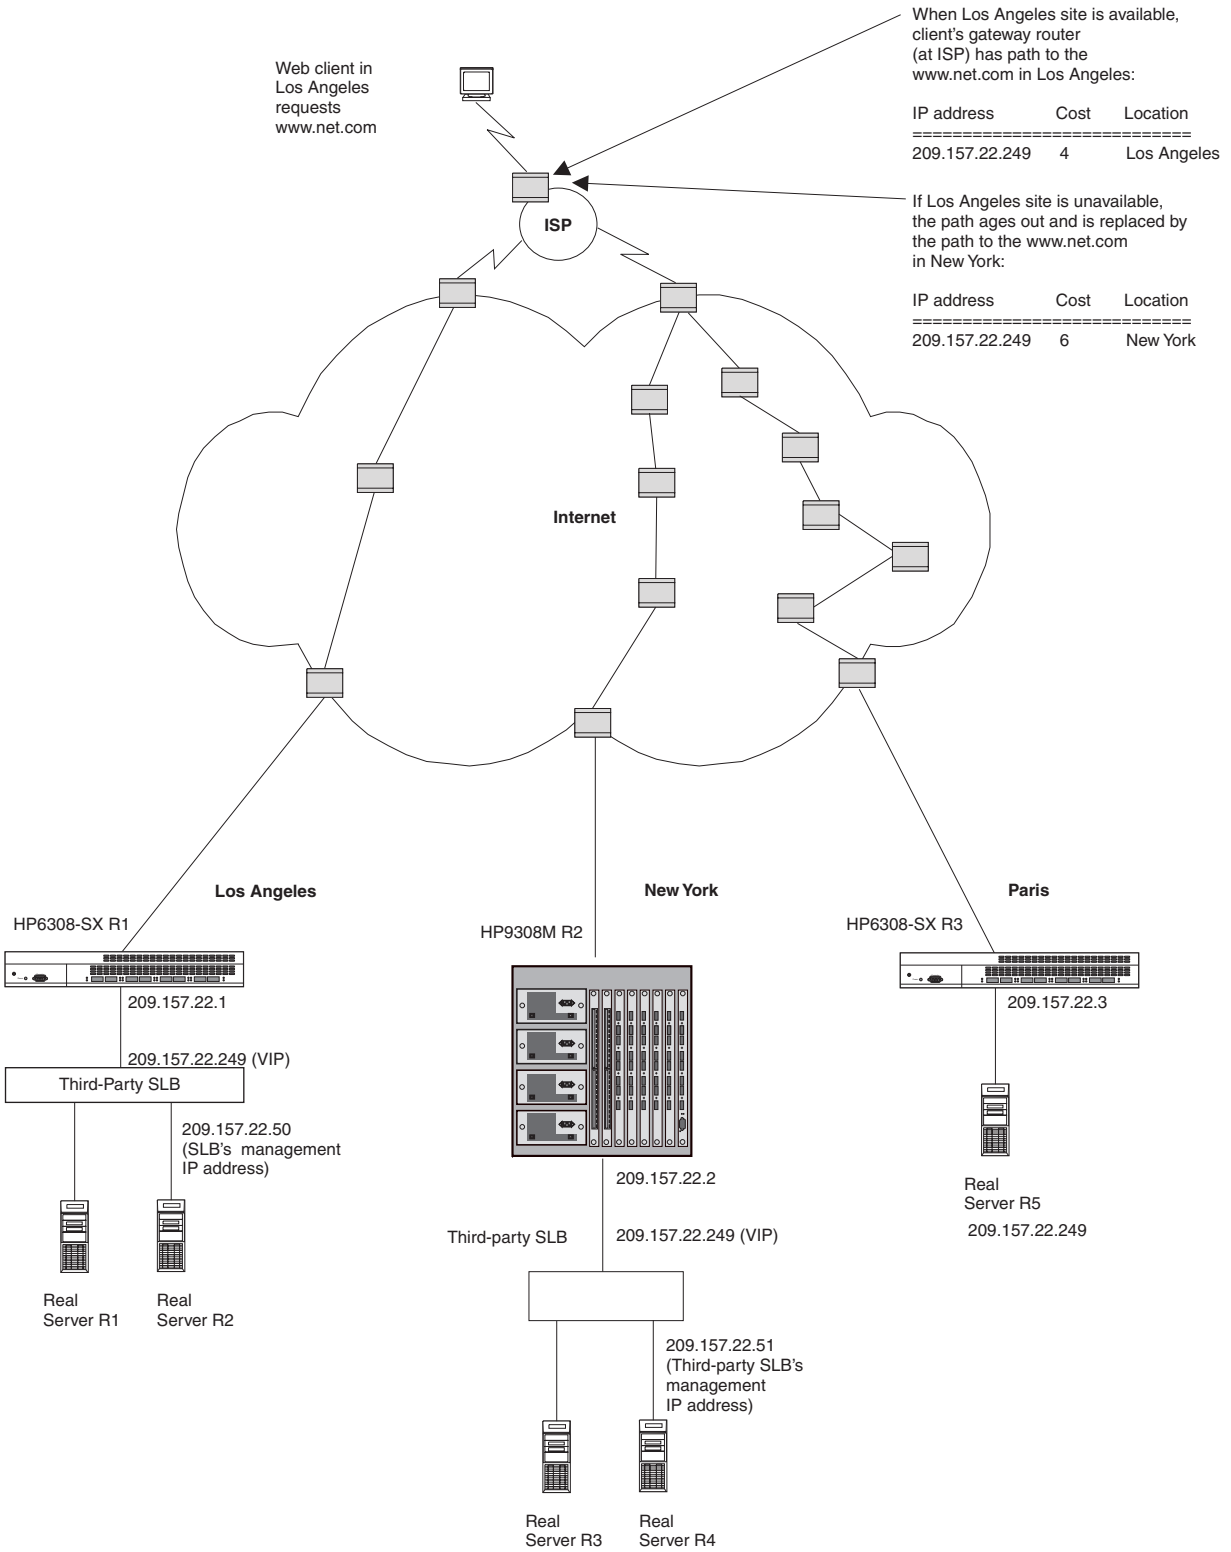


Figure 14.1 Route health injection configuration

When the web browser sends its TCP SYN request (to initiate the HTTP session with the web host), the gateway router used by the client's computer looks in its routing table for the route to the requested IP address. The router may receive multiple paths, in which case the router typically chooses the path with the lowest cost (usually the number of router hops to the host) to place in the routing table. The paths can all go to the same host or to different hosts. In the case of globally-distributed SLB, the paths go to different hosts. The shortest path takes the client to the gateway router attached to the SLB or the directly-attached server that is closest to the client. Thus, when a client on the West coast requests the web site, the client's gateway sends the request to the SLB in Los Angeles. A client in London would instead be directed to the directly-attached server in Paris.

The router's behavior works well when all the real servers are available. However, suppose the real servers attached to the SLB in Los Angeles become unavailable. This results in the VIP on that SLB becoming unavailable.

In a globally-distributed SLB configuration, a client can still reach the desired VIP (web site) if the client's gateway router receives a path to another site that contains the VIP the client is trying to reach. However, gateway routers typically advertise network routes rather than host routes. As a result, even if the VIP (web site) is unavailable, the gateway router still advertises the network to which the VIP belongs. Consequently, a client's gateway router can still have a path to the unavailable server, in which case the client does not receive the requested web page.

By configuring the routing switches attached to the SLBs or real servers that contain the web site to check the health of the web site (HTTP application), you can ensure that the routing switches advertise paths only for web site locations that are available:

- If the web site passes the health check, the routing switch advertises a host route to the web site's IP address.
- If the web site fails the health check, the routing switch removes the host route. The route is no longer advertised and ages out of the routing tables in clients' gateway routers.

As a result, those paths to the web site's IP address that are no longer available age out of the routing tables on gateway routers while the paths that are still available remain in the routing tables. When a client uses its gateway router to reach the web site, the gateway's path to the site's IP address is usually the one with the lowest cost. In Figure 14.1, when the site at Los Angeles is available, the client's gateway uses the path to Los Angeles as the route to IP address 209.157.22.249. However, if the IP address at the Los Angeles site becomes unavailable and thus fails its health check, the HP 6308M-SX routing switch at the Los Angeles site removes the static host route for 209.157.22.249 from its route table. The path on the client's gateway ages out and is replaced by the next valid path with the lowest cost, in this case the path to 209.157.22.249 at the New York site.

HTTP Health Check Algorithm

When you configure a routing switch to periodically check the health of the HTTP port on a web server, the routing switch does one of the following based on the result of the health check. The health check algorithm applies regardless of whether the web server is directly attached to the routing switch (or attached through Layer 2 switches) or is attached to an SLB that is load balancing the IP address among multiple servers.

- If the health check is successful, the routing switch places a static host route in its route table for the web site's IP address. When the routing switch sends a routing advertisement, the host route is included. The client's gateway router will receive this host route as one of the paths to the IP address.
- If the health check is not successful, the routing switch removes the static host route (if present) for the IP address. As a result, the route ages out of the routing tables on other routers. After the removed route ages out of the routing table on the client's gateway router, the router accepts another path to the IP address.

You can configure a separate HTTP health check for each web site IP address. The health check consists of a standard TCP connection followed by a standard request for an HTTP page on the IP address. If the HTTP page responds with an acceptable HTTP status code, the IP address passes the health check, at which point the routing switch leaves the static host route to the IP address in the route table or adds the route if it is not present.

By default, the HTTP health check is disabled. Once you enable the health check, the routing switch sends the health check every five seconds by default. The default health check consists of a HEAD request for the default home page "1.0". If the web site does not respond to a health check, the routing switch resends the health check up to two more times by default before determining that the web site is no longer available and removing the static host route for the web site.

All the health check parameters are configurable. See "CLI Syntax" on page 14-4.

Configuration Considerations

- The routing switch and the SLB or real server must be in the same IP sub-net.
- Place the management station for the SLB on a different sub-net than the one that contains the web site (HTTP application) whose health you are checking. If the web site and the management station are on the same sub-net, the **ip dont_advertise** command (see “CLI Syntax” on page 14-4) will prevent you from reaching the SLB through the management station.
- You cannot use the same routing switch port for OSPF and for the health check. If the port already contains configuration information for one of these features, you cannot configure the other feature unless you first remove the configuration information for the first feature.

CLI Syntax

Use the following commands to configure the health check parameters on a routing switch.

Global CONFIG Level

Use the following command at the global CONFIG level to identify the VIP that has the HTTP port the routing switch is checking.

Syntax: server real <name> <vip>

The <name> parameter identifies the SLB or real server. This value does not need to match a value on the SLB or real server. The value simply identifies the SLB or real server uniquely on the routing switch.

The <vip> parameter is the IP address of the web site. If the web server is directly attached to the routing switch, this is the IP address of the web server. If the web server is attached to an SLB, the VIP is the virtual IP address configured on the SLB for the web site.

Use the following commands to change the interval and retry values for the HTTP health check. When you press Enter after the first command, the CLI changes to the TCP/UDP port configuration level for port 80.

Syntax: server port 80

Syntax: tcp keepalive <interval> <retries>

The <interval> parameter specifies the number of seconds between health checks sent by the routing switch. You can specify a number from 2 – 60 seconds. The default is 5 seconds.

The <retries> parameter specifies how many times the routing switch will resend a health check if the web site does not respond. You can specify from 1 – 5 retries. The default is 2.

Real Server Level

After you enter the **server real...** command shown above, the CLI changes to the Real Server level.

The following command enables the HTTP health check for the web site. The health check is disabled by default.

Syntax: port http keepalive

The following command is optional and changes the default method and URL for the health check. By default, the routing switch sends a HEAD request for the default homepage, “1.0”. The slash in the URL is optional; the routing switch inserts the slash for you if you leave it out.

Syntax: port http url “[GET | HEAD] [/<URL-page-name>”

The following command changes the HTTP status codes that the routing switch accepts as valid responses to a health check. The default status code range for HTTP health checks in SLB configurations is 200 – 299. You can specify up to four discrete ranges. To specify a single message code for a range, enter the code twice. For example to specify 200 only, enter the following command: **port http status_code 200 200**.

Syntax: port http status_code <range> [<range> [<range> [<range>]]]

Interface Level

The following commands configure an IP sub-net address that is in the same sub-net as the web site's IP address. Enter these commands on the interface that connects the routing switch to the real server or to the SLB that is load balancing for the IP address.

The **ip dont_advertise** command configures the routing switch to block advertisement of the host route for the interface. If you do not block the network route, the routing switch will still advertise a network route to the network containing the web site even if the web site itself is unavailable. After you enter the **ip dont_advertise** command, the routing switch advertises only a host route to the IP address. Thus, if the web site fails the HTTP health check, the routing switch removes the static host route for the web site's IP address and also does not advertise a network route for the network containing the IP address.

Syntax: ip address <ip-addr> <ip-mask> [secondary]

Or

Syntax: ip address <ip-addr>/<mask-bits>

Syntax: ip dont_advertise <ip-addr> <ip-mask>

Or

Syntax: ip dont_advertise <ip-addr>/<mask-bits>

Configuring the HTTP Health Check on the Routing Switch

To configure a routing switch to perform the HTTP health check for a web site and to manage a static host route for the IP address, do the following:

- Identify the web site's IP address on the routing switch.
- Enable the HTTP keepalive (health check).
- Optionally modify the health-check keepalive interval and retries.
- Optionally modify site-specific health check parameters (the URL requested by the health check and the HTTP status codes that the routing switch will accept as a normal response).
- Configure the port that connects the routing switch to the HTTP application (SLB or real server) to not advertise the network route for the IP sub-net the SLB or real server and the port are on.

For example, to configure routing switches for the configuration shown in Figure 14.1, enter the following CLI commands.

CLI Commands for 6308M-SX R1

To configure the health check on 6308M-SX R1, enter the following commands:

```
HP6308-R1(config) server real S1 209.157.22.249
HP6308-R1(config-rs-S1) port http keepalive
HP6308-R1(config-rs-S1) interface ethernet 6
HP6308-R1(config-if-6) ip address 209.157.22.1/24
HP6308-R1(config-if-6) ip dont_advertise 209.157.22.1/24
HP6308-R1(config-if-6) write memory
```

Syntax: server real <name> <ip-addr>

Syntax: port http keepalive

Syntax: ip dont_advertise <ip-addr> <ip-mask>

Or

Syntax: ip dont_advertise <ip-addr>/<mask-bits>

The **server real** command in this example configures the HP 6308M-SX to send an HTTP health check to the HTTP port on IP address 209.157.22.249. When you press Enter after this command, the CLI changes to the Real Server level of the CLI. This level allows you to configure health check parameters for the HTTP port on the IP address.

The **port http keepalive** command in this example is entered at the Real Server level and enables the HTTP health check. The health check is disabled by default, so you must enter this command. You can enter additional commands at this level to modify the health check parameters. These commands are shown in the examples for 9308M R2 and 6308M-SX R3.

The **ip address** command adds an IP interface for the connection to the IP address. This interface must be in the same sub-net as the IP address.

The **ip dont_advertise** command configures the routing switch to block advertisement of the network route for this IP sub-net address. This command ensures that the routing switch advertises only the host route to the IP address. If the routing switch advertises the network route to the sub-net containing the IP address, then even if the routing switch removes the host route from its routing table, the routing switch will still advertise the network route to the IP address (and thus to the web server), defeating the failover capability of globally-distributed SLB.

CLI Commands for 9308M R2

The following commands configure 9308M R2 for the configuration shown in Figure 2.

```
HP9300-R2(config) server real S2 209.157.22.249
HP9300-R2(config-rs-S2) port http keepalive
HP9300-R2(config-rs-S2) port http url "/sales.html"
HP9300-R2(config-rs-S2) port http status_code 200 199
HP9300-R2(config-rs-S2) interface ethernet 1/3
HP9300-R2(config-if-1/3) ip address 209.157.22.2/24
HP9300-R2(config-if-1/3) ip dont_advertise 209.157.22.2/24
HP9300-R2(config-if-1/3) write memory
```

Syntax: port http url “[GET | HEAD] [/]<URL-page-name>”

Syntax: port http status_code <range> [<range> [<range> [<range>]]]

The **port http url** command changes the URL that the routing switch sends as part of the health check. By default, the routing switch sends an HTTP HEAD request for the default page (“1.0”). If you enter a URL, the health check instead requests that URL. The slash (/) is an optional parameter. If you do not set the GET or HEAD parameter, and the slash is not in the configured URL page, then the routing switch automatically inserts a slash before retrieving the URL page.

In addition to specifying another URL, you can change the method to GET. Changing the method does not affect the health check from the routing switch’s standpoint. You can use either method.

The **port http status_code** command in this example changes the range of HTTP status codes the routing switch accepts as normal (healthy) replies to a health check.

CLI Commands for 6308M-SX R3

The following commands configure 6308M-SX R3 for the configuration shown in Figure 14.1. This example includes the commands for modifying the HTTP health check interval and retry values.

```
HP6308-R3(config) server port 80
HP6308-R3(config-port-80) tcp keepalive 10 3
HP6308-R3(config-port-80) server real S3 209.157.22.249
HP6308-R3(config-rs-S2) port http keepalive
HP6308-R3(config-rs-S2) port http url "/marketing.html"
HP6308-R3(config-rs-S2) interface ethernet 9
HP6308-R3(config-if-9) ip address 209.157.22.3/24
HP6308-R3(config-if-9) ip dont_advertise 209.157.22.3/24
HP6308-R3(config-if-9) write memory
```

Syntax: server port 80

Syntax: tcp keepalive <interval> <retries>

The <interval> parameter specifies the number of seconds between health checks sent by the routing switch. You can specify a number from 2 – 60 seconds. The default is 5 seconds.

The <retries> parameter specifies how many times the routing switch will resend a health check if the web site does not respond. You can specify from 1 – 5 retries. The default is 2.

Displaying Server and Application Port Information

You can use the CLI to display the following types of information:

- Server (virtual IP address) information
- Application port information

Displaying Server Information

To display information about the server virtual IP addresses (VIPs) you have configured, enter a command such as the following at any level of the CLI:

```
HP9300-R2# show server real RS2

Real Servers Info

Server State - 1:enabled, 2:failed, 3:test, 4:suspect, 5:grace_dn, 6:active
Name:RS2          IP: 209.157.23.60:4      State:1
```

Syntax: show server real <name>

This display shows the following information.

Table 14.1: Real Server Information

This Field...	Displays...
Server State	The possible values for the server state. The state of each real server is shown by the State field. See below.
Name	The name of the real server. This is the name you assigned to the server when you configured it on the SLB.

Table 14.1: Real Server Information (Continued)

This Field...	Displays...
IP	The IP address of the real server. If you configured a host range of VIPs on the server, the number following the IP address (after the colon) is the number of hosts on the server.
State	The state of the real server. The state can be one of the states listed by "Server State" at the top of the display.

Displaying Keepalive Information

To display the keepalive parameters in effect for the application ports on the servers, enter the following command at any level of the CLI:

Syntax: show server keepalive-port