

Chapter 19

Redirecting Subscriber Traffic

This chapter describes the redirect server and contains the following sections:

- Overview of Traffic Redirection on page 345
- Proxy Request Management on page 345
- Redirect Server Redundancy on page 347
- Before You Configure Redundancy for the Redirect Server on page 348
- Protection Against Denial-of-Service Attacks on page 348

Overview of Traffic Redirection

The redirect server is part of a captive portal system that redirects subscribers' Web requests to a captive portal page. You can use a captive portal page as the initial page a subscriber sees after logging in to a subscriber session and as a page used to receive and manage HTTP requests to unauthorized Web resources.

If you run the SRC software on Solaris platforms, a captive portal uses an instance of the redirect server installed on a host in the same network as a JUNOSe router. Your network configuration must not have any routers between the JUNOSe router and the redirect server. An intermediate router would look at the destination address that is still present in the packets and would route the packets there rather than to the SAE. One way to overcome this limitation is to set up a tunnel between the JUNOSe router and the redirect server.

Proxy Request Management

The redirect server examines requested paths and detects proxy HTTP requests by the proxy prefix "< scheme >:" followed by the address of the requested host. If the requested URL is served by the captive portal server:

1. The redirect server opens a TCP connection to the captive portal and forwards the request for the URL. The redirect server adds to the request an X-Forwarded-For header that specifies the IP address of the client.

2. The captive portal server inspects the incoming request for the X-Forwarded-For header for the IP address. The captive portal server uses this address instead of the source IP address to determine the originator of the request.
3. If the captive portal authorizes the client and activates a service that enables a direct connection between the client and the proxy, the redirect server then sends the returned data to the subscriber's Web browser.

or

If the requested URL is not served by the captive portal server, the redirect server opens a TCP port (8800 by default) and sends the type of response configured to a subscriber's browser in response to a captured request:

- HTTP 200 OK response with an HTML document that includes the < HTTP-Equiv = "Refresh" > header (default)
- HTTP 302 Found response to a subscriber's browser in response to a captured request

The subscriber browser follows the redirect request, and the proxied request is served by the redirect server again, which opens a connection to the captive portal.

Support for HTTP proxy requests requires the following:

- A local HTTP proxy server that can handle the traffic from all clients configured with a proxy.
- A location for the local HTTP proxy server that is one IP hop from each access router.
- A proxy service that the captive portal server can activate to send proxy requests to the local HTTP proxy server when the portal server authorizes proxy clients.
- A proxy service activation policy that includes a next-hop policy that points to the local HTTP proxy server, and a classifier that matches the client's IP address and the address of the proxy server configured on the client.

Services that the client accesses through the proxy server, such as HTTP and FTP, cannot be activated based on destination address.

You must redirect all ports to the redirect server because you cannot know which ports are configured on the client for the proxy. Consequently, the redirect server receives non-HTTP requests as well as HTTP requests. The non-HTTP requests generate error log entries. To reduce overhead, HTTP error messages are logged as system log debug messages.

HTTP Proxy and DNS

Make sure that your network includes a domain name service (DNS) server to resolve unknown names to a fixed IP address. A DNS server is required because proxy servers can be configured with DNS names in private domains that are not valid in the public environment. You can use the DNS server included with the redirect server, or another DNS server on your network.

The DNS server can be configured on a client with DHCP. Alternatively, the service provider can set up a transparent DNS proxy by configuring a next-hop policy on the JUNOS router for UDP and TCP port 53 traffic. The policy redirects traffic on these two ports to the redirect server's DNS server.

Because proxy addresses must be resolved even if general access to the Internet is enabled, the DNS server must continue to resolve all client requests for proxy clients. Nonproxy clients can use their regular DNS server after the initial service has been activated.

The redirect server's DNS server either forwards the request to a set of configured DNS servers or resolves the request by using the root domain name server. If a request for an IPv4 address cannot be resolved and the request results in an NXDOMAIN error, the DNS server returns a configurable IP address. The redirect server returns an error message to the clients for any other type of request that cannot be resolved.


Redirect Server Redundancy

You can configure the redirect server to provide redundancy to help ensure that a redirect server is always available. You install the redirect server software on two different hosts; then you configure one redirect server as the primary redirect server, and the other as the redundant redirect server. The active and redundant redirect servers regularly poll each other to confirm each other's availability. If the primary redirect server becomes unavailable, the redundant server assumes the active role.

When a redirect server assumes the primary role, it configures on the router a static route from the virtual IP address to the server's real IP address. Clients send requests to the virtual IP address, and the router automatically sends the request to the active redirect server through a static route. The virtual IP address is used only in the static route configured on the router and the next-hop policy installed by SAE. End users do not see the virtual IP address.

Figure 28 shows a configuration in which two redirect servers use the same virtual IP address, 192.168.254.1.

Figure 28: Failover of Redirect Server



Before You Configure Redundancy for the Redirect Server

If you plan to use a redundant configuration for the redirect server, ensure that:

- The virtual IP address to be used is also the next-hop address for policies that capture web traffic and send it to the redirect server.
- The redirect server has SNMP write access to the virtual routers connected to it. Each VR must have at least a write community configured. (The static route from the virtual IP address to the server's real IP address is installed on the router through SNMP.)
- If additional access controls are enabled on the JUNOS router, the hosts on which the redirect server runs must be included.

Protection Against Denial-of-Service Attacks

The redirect server incorporates a number of properties to protect against denial-of-service attacks. The following list shows the default values set for these properties:

- The redirect server can serve no more than 12,000 requests per minute, with a burst of 18,000 requests.
- The redirect server can serve no more than 25 requests per client per minute, with a burst of 50 requests.
- Incoming requests can be no larger than 4 KB.
- Incoming requests have a time limit of 2 seconds.

You can change the values for any of these properties.