

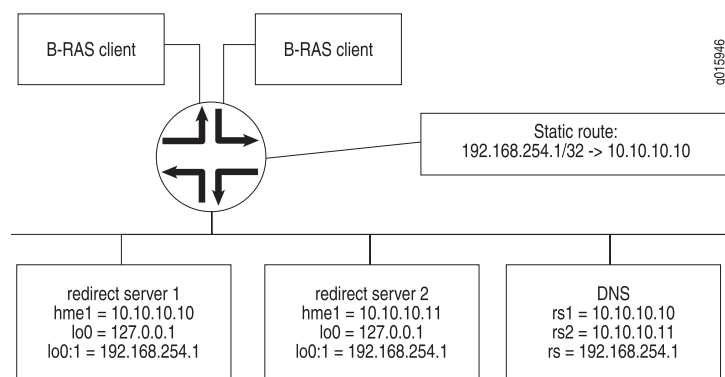
Redirect Server Redundancy

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 1 on page 1 shows a configuration in which two redirect servers use the same virtual IP address, 192.168.254.1.

Figure 1: Failover of a Redirect Server



- Related Topics**
- Overview of Traffic Redirection
 - Before You Configure Redundancy for a Redirect Server
 - Configuring a Redundant Redirect Server (SRC CLI)
 - Configuring a Redundant Redirect Server with the C-Web Interface

