Chapter 12
Configuring BGP4

This chapter provides details on how to configure ***Border Gateway Protocol version 4 (BGP4)*** on the HP ProCurve 9304M, 9308M, and 6308M-SX routing switches using the CLI and the Web management interface.

BGP4 is described in RFC 1771. The BGP4 implementation on the 9304M, 9308M, and 6308M-SX routing switches fully complies with RFC 1771. The implementation also includes RFC 1745 (OSPF Interactions) and RFC 1997 (BGP Communities Attributes).

To display BGP4 configuration information and statistics, see “Displaying BGP4 Information” on page 12-43.

This chapter shows the commands you need in order to configure the routing switches for BGP4. For a detailed list of all CLI commands, including syntax and possible values, see “Command Line Interface Commands” on page B-1.

NOTE: Your routing switch must have 32MB or higher to run BGP4.

NOTE: 9304M and 9308M routing switches that use Redundant Management modules can contain a maximum of 80000 IP routes by default. The 6308M-SX and chassis devices that use other management modules can contain a maximum of 10000 IP routes by default. If you need to increase the capacity of the IP route table for BGP4, see “Modifying System Parameter Default Settings” on page 8-69.

Overview of BGP4

BGP4 is the standard Exterior Gateway Protocol (EGP) used on the Internet to route traffic between ***Autonomous Systems (AS)*** and to maintain loop-free routing. An autonomous system is a collection of networks that share the same routing and administration characteristics. For example, a corporate intranet consisting of several networks under common administrative control might be considered an AS. The networks in an AS can but do not need to run the same routing protocol to be in the same AS, nor do they need to be geographically close.

Routers within an AS can use different Interior Gateway Protocols (IGPs) such as RIP and OSPF to communicate with one another. However, for routers in different ASs to communicate, they need to use an EGP. BGP4 is the standard EGP used by Internet routers and therefore is the EGP implemented on the 9304M, 9308M, and 6308M-SX routing switches.

Figure 12.1 shows a simple example of two BGP4 ASs. Each AS contains three BGP4 routers. All of the BGP4 routers within an AS communicate using IBGP. BGP4 routers communicate with other ASs using EBGP. Notice that each of the routers also is running an IGP. The routers in AS1 are running OSPF and the routers in AS2 are running RIP. The 9304M, 9308M, and 6308M-SX routing switches can be configured to redistribute routes among BGP4, RIP, and OSPF. They also can redistribute static routes.

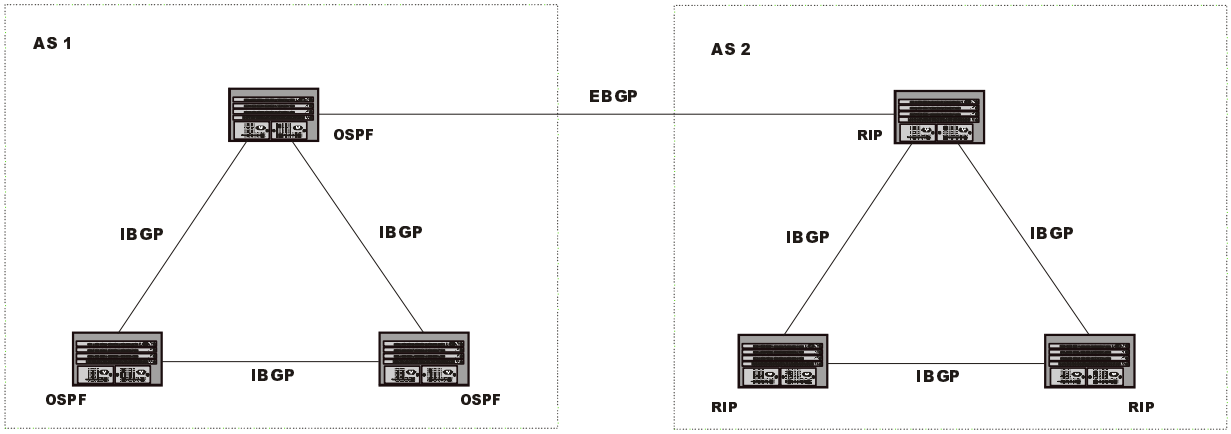

Figure 12.1 Example BGP4 ASs

Relationship Between the BGP4 Route Table and the IP Route Table

The routing switch's BGP4 route table can have multiple routes to the same destination, which are learned from different BGP4 neighbors. A BGP4 neighbor is another router that also is running BGP4. BGP4 neighbors communicate using Transmission Control Protocol (TCP) port 179 for BGP communication. When you configure a routing switch for BGP4, one of the configuration tasks you perform is to identify the routing switch's BGP4 neighbors.

Although a routing switch's BGP4 route table can have multiple routes to the same destination, the BGP4 protocol evaluates the routes and chooses only one of the routes to send to the IP route table. The route that BGP4 chooses and sends to the IP route table is the ***preferred route*** and will be used by the routing switch. If the preferred route goes down, BGP4 updates the route information in the IP route table with a new BGP4 preferred route.

A BGP4 route consists of the following information:

- Network number (prefix) – A value comprised of the network mask bits and an IP address (<IP address>/<mask bits>); for example, 192.215.129.0/18 indicates a network mask of 18 bits applied to the IP address 192.215.129.0. When a BGP4 routing switch advertises a route to one of its neighbors, the route is expressed in this format.
- AS-path – A list of the other ASs through which a route passes. BGP4 routers can use the AS-path to detect and eliminate routing loops. For example, if a route received by a BGP4 router contains the AS that the router is in, the router does not add the route to its own BGP4 table. (The BGP4 RFCs refer to the AS-path as "AS_PATH".)
- Additional path attributes – A list of additional parameters that describe the route. The route origin and next hop are examples of these additional path attributes.

After a routing switch successfully negotiates a BGP4 session with a neighbor (a BGP4 peer), the routing switch exchanges complete BGP4 route tables with the neighbor. After this initial exchange, the routing switch and all other RFC 1771-compliant BGP4 routers send UPDATE messages to inform neighbors of new, changed, or no longer feasible routes. BGP4 routers do not send regular updates. However, if configured to do so, a BGP4 router does regularly send KEEPALIVE messages to its peers to maintain BGP4 sessions with them if the router does not have any route information to send in an UPDATE message. See "BGP4 Message Types" on page 12-4 for information about BGP4 messages.

How BGP4 Selects a Path for a Route

When multiple paths for the same route are known to a BGP4 router, the router uses an algorithm to weigh the paths and determine the optimal path for the route. The optimal path depends on various parameters including the following. You can modify some of these parameters. (See “Optional Configuration Tasks” on page 12-16.)

- Weight – A value that the BGP4 router associates with a specific BGP4 neighbor. For example, if the routing switch receives routes to the same destination from two BGP4 neighbors, the routing switch prefers the route from the neighbor with the larger weight.
- Local preference – An attribute that indicates a degree of preference for a route relative to other routes in the local AS.
- AS-path length – The number of ASs through which the route must pass to reach the destination. The AS-path is a sequential list of the AS numbers through which the route information has passed to reach the BGP4 routing switch.
- Origin – The source of the route information. The origin can be IGP, EGP, or INCOMPLETE. IGP is preferred over EGP and both are preferred over INCOMPLETE.
- Multi-Exit Discriminator (MED) – A value associated with routes that have multiple paths through the same AS. In BGP4, a route’s MED is equivalent to its “metric”.
- Closest IBGP neighbor – The closest internal path to the destination within the local AS.

The 9304M, 9308M, and 6308M-SX routing switches use the following algorithm to choose the optimal path for a BGP4 route. The algorithm uses the parameters listed above.

1. Is the next hop in the route accessible? If not, ignore the route.
2. Use the path with the largest weight.
3. If the weights are the same, prefer the route with the largest local preference.
4. If the routes have the same local preference, prefer the route that was originated locally (by this BGP4 routing switch).
5. If the local preferences are the same and the routes were originated locally, prefer the route with the shortest AS-path.
6. If the AS-path lengths are the same, prefer the route with the lowest origin type. From low to high, route origin types are valued as follows:
   - IGP is lowest.
   - EGP is higher than IGP but lower than INCOMPLETE.
   - INCOMPLETE is highest.
7. If the routes have the same origin type, prefer the route with the lowest MED.
8. If the routes have the same MED, prefer routes in the following order:
   - Routes received through EBGP
   - Routes received through IBGP
9. If all the comparisons above are equal, prefer the route that can be reached using the closest IGP neighbor. This is the closest internal path inside the AS to reach the destination.
10. If the internal paths also are the same, prefer the route that comes from the BGP4 router with the lowest router ID.

BGP4 Message Types

BGP4 routers communicate with their neighbors (other BGP4 routers) using the following types of messages:

- OPEN
- UPDATE
- KEEPALIVE
- NOTIFICATION

OPEN Message

After a BGP4 router establishes a TCP connection with a neighboring BGP4 router, the routers exchange OPEN messages. An OPEN message indicates the following:

- BGP version – Indicates the version of the protocol that is in use on the router. BGP version 4 supports Classless Interdomain Routing (CIDR) and is the version most widely used in the Internet. Version 4 also is the only version supported on the 9304M, 9308M, and 6308M-SX routing switches.
- AS number – A two-byte number that identifies the AS to which the BGP4 router belongs.
- Hold Time – The number of seconds a BGP4 router will wait for an UPDATE or KEEPALIVE message (described below) from a BGP4 neighbor before assuming that the neighbor is dead. BGP4 routers exchange UPDATE and KEEPALIVE messages to update route information and maintain communication. If BGP4 neighbors are using different Hold Times, the lowest Hold Time is used by the neighbors. If the Hold Time expires, the BGP4 router closes its TCP connection to the neighbor and clears any information it has learned from the neighbor and cached.

  You can configure the Hold Time to be 0, in which case a BGP4 router will consider its neighbors to always be up. For directly-attached neighbors, the router is configured by default to immediately close the TCP connection to the neighbor and clear entries learned from an EBGP neighbor if the interface to that neighbor goes down. This capability is provided by the fast external fallover feature, which is enabled by default.
- BGP Identifier – The router ID. The BGP Identifier (router ID) identifies the BGP4 router to other BGP4 routers. The 9304M, 9308M, and 6308M-SX routing switches use the same router ID for OSPF and BGP4. If you do not set a router ID, the software uses the lowest IP address configured on the router.
- Parameter list – An optional list of additional parameters used in peer negotiation with BGP4 neighbors.

UPDATE Message

After BGP4 neighbors establish a BGP4 connection over TCP and exchange their BGP4 routing tables, they do not send periodic routing updates. Instead, a BGP4 neighbor sends an update to its neighbor when it has a new route to advertise or routes have changed or become unfeasible. An UPDATE message can contain the following information:

- Network Layer Reachability Information (NLRI) – The mechanism by which BGP4 supports Classless Interdomain Routing (CIDR). An NLRI entry consists of an IP prefix that indicates a network being advertised by the UPDATE message. The prefix consists of an IP network number and the length of the network portion of the number. For example, an UPDATE message with the NLRI entry 192.215.129.0/18 indicates a route to IP network 192.215.129.0 with network mask 255.255.192.0. The binary equivalent of this mask is 18 consecutive one bits, thus “18” in the NLRI entry.
- Path attributes – Parameters that indicate route-specific information such as path information, route preference, next hop values, and aggregation information. BGP4 uses the path attributes to make filtering and routing decisions.
- Unreachable routes – A list of routes that have been in the sending router’s BGP4 table but are no longer feasible. The UPDATE message lists unreachable routes in the same format as new routes: <IP address>/<CIDR prefix>.

KEEPALIVE Message

BGP4 routers do not regularly exchange UPDATE messages to maintain the BGP4 sessions. For example, if a 9308M configured to perform BGP4 routing has already sent the latest route information to its peers in UPDATE messages, the router does not send more UPDATE messages. Instead, BGP4 routers send KEEPALIVE messages to maintain the BGP4 sessions. KEEPALIVE messages are 19 bytes long and consist only of a message header; they contain no routing data.

BGP4 routers send KEEPALIVE messages at a regular interval, the Keep Alive Time. The default Keep Alive Time on routing switches is 60 seconds.

A parameter related to the Keep Alive Time is the Hold Time. A BGP4 router's Hold Time determines how many seconds the router will wait for a KEEPALIVE or UPDATE message from a BGP4 neighbor before deciding that the neighbor is dead. The Hold Time is negotiated when BGP4 routers exchange OPEN messages; the lower Hold Time is then used by both neighbors. For example, if BGP4 Router A sends a Hold Time of 5 seconds and BGP4 Router B sends a Hold Time of 4 seconds, both routers use 4 seconds as the Hold Time for their BGP4 session. The default Hold Time is 180 seconds. Generally, the Hold Time is configured to three times the value of the Keep Alive Time.

If the Hold Time is 0, a BGP4 router assumes that its neighbor is alive regardless of how many seconds pass between receipt of UPDATE or KEEPALIVE messages.

NOTIFICATION Message

When you close the router's BGP4 session with a neighbor, or the router detects an error in a message received from the neighbor, or an error occurs on the router, the router sends a NOTIFICATION message to the neighbor. No further communication takes place between the BGP4 router that sent the NOTIFICATION and the neighbor(s) that received the NOTIFICATION.

Basic Configuration and Activation for BGP4

BGP4 is disabled by default. To enable BGP4 and place your routing switch into service as a BGP4 router, you must perform at least the following steps:

1. Enable the BGP4 protocol.
2. Set the local AS number (unless you want to use the default AS number "1").
3. Add each BGP4 neighbor (peer BGP4 router) and identify the AS the neighbor is in. Each routing switch can have up to ten BGP4 neighbors. The default number of neighbors allowed is three.
4. Save the BGP4 configuration information to the system configuration file.

NOTE: By default, the router ID is the lowest IP address configured on the routing switch. If you want to assign a different router ID, use the **ip router-id** command. (See "ip router-id" on page B-102.)

The default local AS number is 1. To change the local AS, see "Setting the Local AS Number" on page 12-12.

USING THE CLI

```
HP9300> enable

HP9300# configure terminal

HP9300(config)# router bgp

HP9300(config-bgp-router)# neighbor <router ID> remote-as <AS number>

HP9300(config-bgp-router)# write mem
```

NOTE: When BGP4 is enabled, you do not need to reset the system. The protocol is activated as soon as you enable it. Moreover, the routing switch begins a BGP4 session with a BGP4 neighbor as soon as you add the neighbor.

USING THE WEB MANAGEMENT INTERFACE

1. Click on the System link to display the System configuration sheet, as shown in Figure 12.2.
2. Select Enable next to BGP.
3. Select Apply to implement the change. BGP4 is activated as soon as you select Apply.

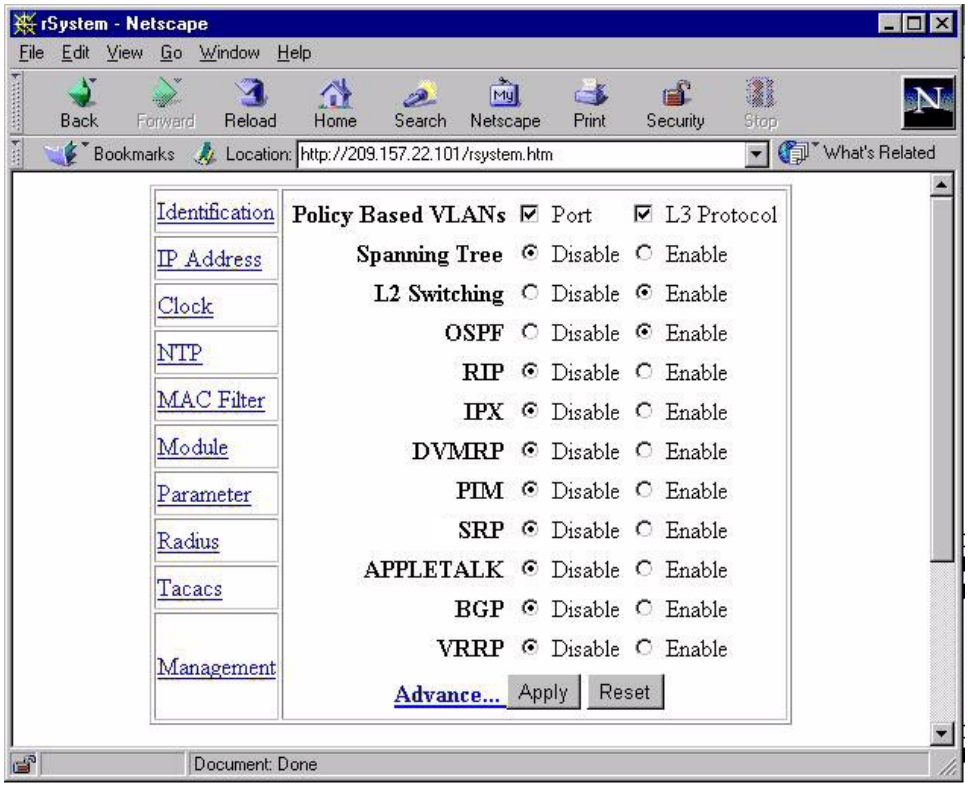

Figure 12.2 System configuration sheet

BGP4 Parameters

You can modify or set the following BGP4 parameters.

- Optional – Define the router ID. (The same router ID also is used by OSPF.)
- Optional – Change the local AS number from 1 (the default).
- Optional – Add a loopback interface for use with IBGP neighbors (neighbors in the same AS).
- Required – Identify BGP4 neighbors.
- Optional – Change the Keep Alive Time and Hold Time.
- Optional – Disable fast external fallover.
- Optional – Change the maximum number of BGP4 neighbors the routing switch can have.
- Optional – Change the maximum number of BGP4 routes the routing switch's BGP4 route table can contain.
- Optional – Change the maximum number of route-attribute entries the routing switch can manage.
- Optional – Specify a list of individual networks in the local AS to be advertised to remote ASs using BGP4.
- Optional – Change the default local preference for routes.

- Optional – Change the default information originate.
- Optional – Change the default MED (metric).
- Optional – Change the default administrative distances for EBGP, IBGP, and local (directly attached) routes.
- Optional – Always compare MEDs (metrics) when choosing a route.
- Optional – Disable synchronization of routes between BGP4 and IGP.
- Optional – Enable auto summary to summarize routes at an IP class boundary (A, B, or C).
- Optional – Aggregate routes in the BGP4 route table into CIDR blocks.
- Optional – Configure the routing switch as a BGP4 router reflector
- Optional – Change the default metric for routes that BGP4 redistributes into RIP or OSPF.
- Optional – Change the parameters for RIP, OSPF, or static routes redistributed into BGP4.
- Optional – Define BGP4 address filters.
- Optional – Define BGP4 AS-path filters.
- Optional – Define BGP4 community filters.
- Optional – Define BGP4 route maps for filtering routes redistributed into RIP and OSPF.

NOTE: When using CLI, you set global level parameters at the BGP CONFIG Level of the CLI. You can reach the BGP CONFIG level by entering **router bgp**… at the global CONFIG level.

NOTE: When using the Web management interface, you set BGP4 global parameters using the BGP4 configuration sheet, shown in Figure 12.3. You can access all other parameters using links on the BGP4 configuration sheet. Select the BGP link from the menu at the bottom or in the lefthand frame of the Web management display.

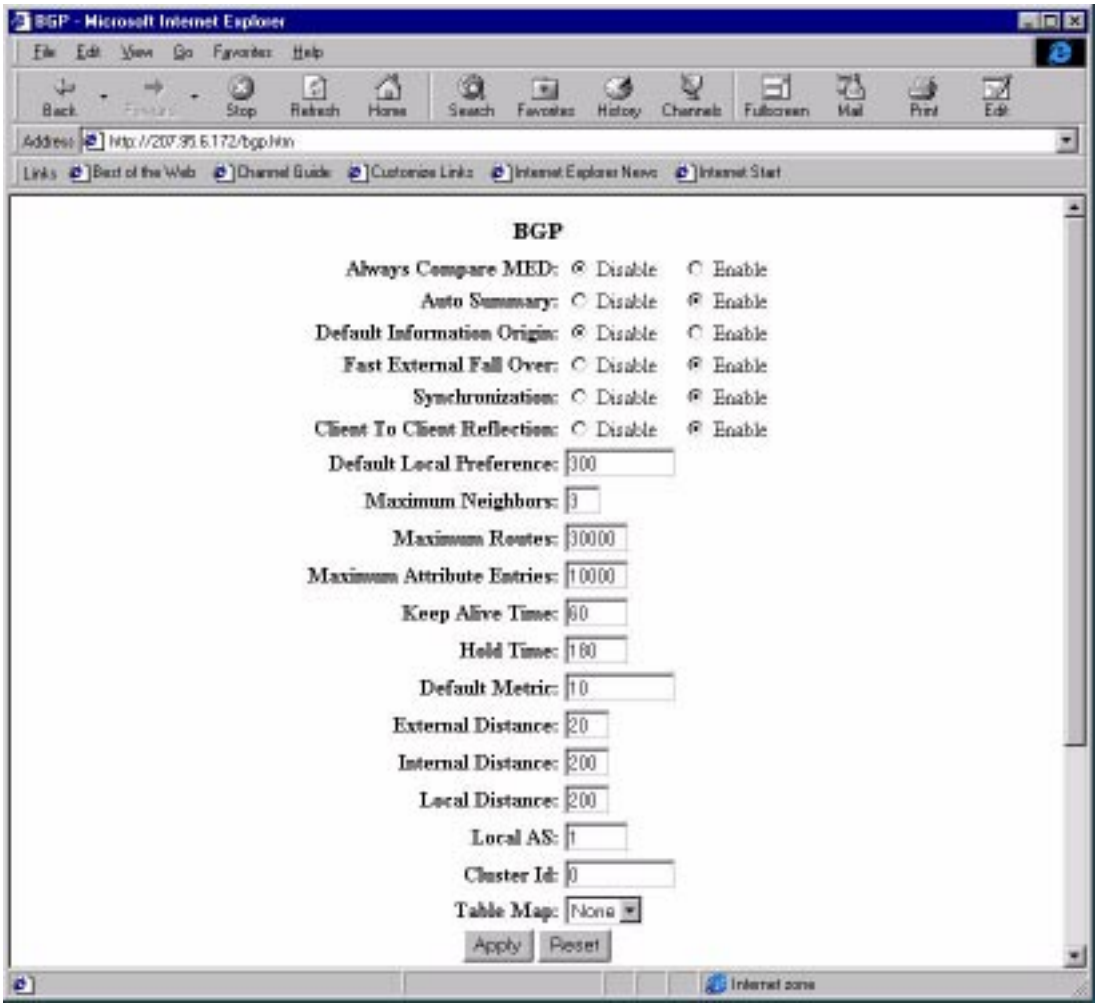

Figure 12.3 BGP configuration sheet

When Parameter Changes Take Effect

Some parameter changes take effect immediately while others do not take full effect until the routing switch's sessions with its neighbors are closed, then restarted. Some parameters do not take effect until the routing switch is rebooted.

Immediately

The following parameter changes take effect immediately:

- Enable or disable BGP.
- Change local AS.
- Add neighbors.
- Enable or disable fast external fallover.
- Specify individual networks that can be advertised.
- Change the default local preference, default information originate, or administrative distance.
- Enable or disable MED (metric) comparison.
- Enable or disable IGP and BGP4 synchronization.

- Disable or re-enable auto summary.
- Change the default metric.
- Disable or re-enable route reflection.

After Resetting Neighbor Sessions

The following parameter changes take effect only after the routing switch's BGP4 sessions are closed, then reopened:

- Change the Hold Time or Keep Alive Time.
- Aggregate routes.
- Add, change, or negate filter tables.
- Add, change, or negate route maps.
- Add, change, or negate redistribution parameters (except changing the default MED; see below).

NOTE: Depending on where filters and route maps are applied, you might also need to disable, then re-enable BGP4. When you disable BGP4 on a routing switch, that routing switch's neighbors clear all the routes they learned from the routing switch. If you want the routing switch to resend its routing table without disabling and re-enabling BGP4, you can use the soft-outbound option when clearing routes learned from a neighbor. See "Closing or Resetting Sessions With Neighbors" on page 12-59.

After Disabling and Re-Enabling Redistribution

The following parameter change takes effect only after you disable and then re-enable redistribution:

- Change the default MED (metric).

After Rebooting or Reloading the Routing Switch

The following parameter changes take effect only after you reboot or reload the routing switch:

- Change the maximum number of BGP4 neighbors, routes, or route-attribute entries the routing switch can have.

Memory Considerations

BGP4 handles a very large number of routes and therefore requires a lot of memory. For example, in a typical configuration with just a single BGP4 neighbor, a BGP4 router may need to be able to hold up to 50,000 routes. Many configurations, especially those involving more than one neighbor, can require the routing switch to hold even more routes.

The 9304M, 9308M, and 6308M-SX routing switches can hold 30,000 routes by default. You can change the route table capacity to a value from 10,000 – 70,000.

NOTE: On a 9304M or 9308M routing switch using Redundant Management modules, the default maximum number of routes is 256,000 and the default maximum number of route-attribute entries is 100,000.

The software reserves memory for BGP4 route tables and other tables such as the IP route table. The memory reserved for the BGP4 route table cannot be used by other parts of the system. Therefore, the system will not operate properly unless all of the features you want to use have adequate memory. In addition, individual ports require memory, so you must ensure that the memory you allocate to route tables and other tables with configurable sizes leaves enough memory for the ports.

- You can configure memory to allow the maximum number of routes only if you also reconfigure memory to allow only one or two neighbors.
- With three or more neighbors, you cannot have the maximum number of routes.

See "Changing the Maximum Number of Neighbors" on page 12-17 and "Changing the Maximum Number of Routes" on page 12-18 for memory configuration procedures.

Configuring BGP4

To begin using BGP4 on the routing switch, follow the steps outlined below:

1. Enable the BGP4 feature on the routing switch.
2. Optionally define the router ID.
3. Optionally change the local AS number from 1 (the default).
4. Identify the routing switch's BGP4 neighbors and the ASs they are in.
5. Optionally change the Keep Alive Time and Hold TIme.
6. Optionally disable fast external fallover.
7. Optionally change the maximum number of BGP4 neighbors, routes, or route-attribute entries the routing switch can have.
8. Optionally specify a list of individual networks in the local AS to be advertised to remote ASs using BGP4.
9. Optionally change the default local preference, default information originate, default MED (metric), or administrative distances. (You change these parameters independently of one another.)
10. Optionally configure the routing switch to always compare MEDs (metrics) when choosing a route.
11. Optionally disable synchronization of routes between BGP4 and IGP.
12. Optionally disable automatic summarization of subnets at the classical IP boundaries (classes A, B, and C).
13. Optionally aggregate routes in the BGP4 route table into CIDR blocks.
14. Optionally configure the routing switch as a BGP4 route reflector.
15. Optionally change the default metric for routes that BGP4 redistributes into RIP or OSPF.
16. Optionally define BGP4 address filters, AS-path filters, or community filters.
17. Optionally define BGP4 route map entries.
18. Save the changes to flash memory.

Basic Configuration Tasks

The following sections describe how to perform the configuration tasks that are required to use BGP4 on the routing switch. You can modify many parameters in addition to the ones described in this section. See "Optional Configuration Tasks" on page 12-16.

Enabling BGP4 on the Routing Switch

When you enable BGP4 on the routing switch, BGP4 is automatically activated. To enable BGP4 on the routing switch, enter the following commands:

USING THE CLI

```
HP9300(config)# router bgp
```

This command places you in the BGP4 router level where you can modify BGP4 global parameters.

USING THE WEB MANAGEMENT INTERFACE

1. Click on the System link to display the System configuration sheet, as shown in Figure 12.2.
2. Select Enable next to BGP.
3. Select Apply to implement the change. BGP4 is activated as soon as you select Apply.

Changing the Router ID

The OSPF and BGP4 protocols use router IDs to identify the routers that are running the protocols. A router ID is a valid, unique IP address and sometimes is an IP address configured on the routing switch. The router ID cannot be an IP address in use by another device. By default, the router ID is the lowest IP address configured on the routing switch. However, you can set the router ID to any valid IP address.

NOTE: The 9304M, 9308M, and 6308M-SX routing switches use the same router ID for both OSPF and BGP4. If the routing switch is already configured for OSPF, you may want to use the router ID that is already in use on the routing switch rather than set a new one. To display the router ID, enter the **show ip** CLI command at any CLI level or select the IP link in the Web management interface.

USING THE CLI

To set the router ID, enter a command such as the following:

```
HP9300(config)# ip router-id 209.157.22.26
```

Syntax: ip router-id <IP-addr>

The <IP-addr> can be any valid, unique IP address.

NOTE: You can specify an IP address used for an interface on the routing switch, but do not specify an IP address in use by another device.

USING THE WEB MANAGEMENT INTERFACE

1. Select the IP link to display the IP configuration sheet, shown in Figure 12.4.
2. Edit the value in the Router ID field to any valid IP address not in use on another router.
3. Select Apply to implement the change.

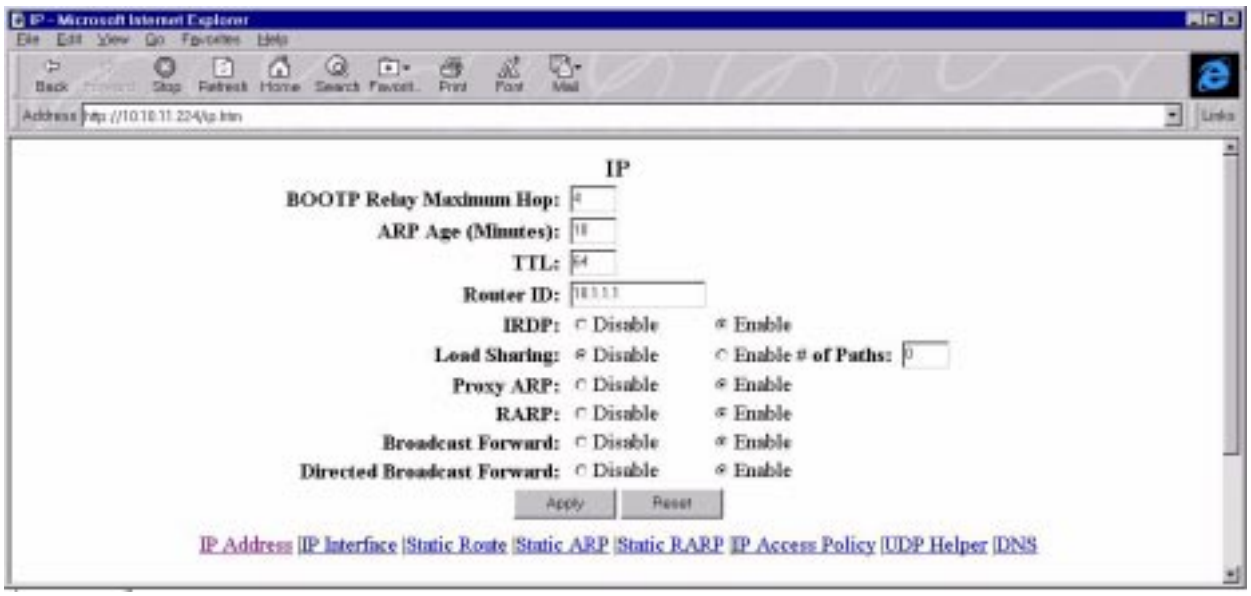

Figure 12.4 IP configuration sheet

Setting the Local AS Number

The local AS number identifies the AS the routing switch is in. The AS number can be from 1 – 65535. The default local AS number is 1. AS numbers 64512 – 65535 are the well-known private BGP4 AS numbers. If you want to set a different local AS value, use either of the following methods.

USING THE CLI

To set the local AS, enter a command such as the following:

```
HP9300(config-bgp-router)# local-as 64512
```

Syntax: local-as <num>

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Enter or edit the number in the Local AS field.
3. Select Apply to implement the change.

Adding a Loopback Interface

You can configure the routing switch to use a loopback interface instead of a specific port to communicate with a BGP neighbor. A loopback interface adds stability to the network by working around route flap problems that can occur due to unstable links between the routing switch and its neighbors.

Loopback interfaces are always up, regardless of the states of physical interfaces. Loopback interfaces are especially useful for IBGP neighbors (neighbors in the same AS) that are multiple hops away from the routing switch. When you configure a BGP neighbor on the routing switch, you can specify whether the routing switch uses the loopback interface to communicate with the neighbor. As long as a path exists between the routing switch and its neighbor, BGP4 information can be exchanged. The BGP4 session is not associated with a specific link but instead is associated with the virtual interfaces.

You can add up to 24 IP addresses to each loopback interface.

NOTE: If you configure the routing switch to use a loopback interface to communicate with a BGP4 neighbor, you also must configure a loopback interface on the neighbor and configure the neighbor to use that loopback interface to communicate with the routing switch.

To add a loopback interface, use one of the following methods.

USING THE CLI

To add a loopback interface, enter commands such as those shown in the following example:

```
HP9300(config-bgp-router)# exit
HP9300(config)# int loopback 1
HP9300(config-lbif-1)# ip address 10.0.0.1/24
```

Syntax: interface loopback <num>

The <num> value can be from 1 – 8.

USING THE WEB MANAGEMENT INTERFACE

1. Select IP.
2. Select Loopback to display the Router Loopback panel.

NOTE: If you have already configured a loopback interface, the configured loopback interfaces are listed. Click Add Loopback.

3. Select the loopback interface number from the Loopback field's pulldown menu. You can select from 1 – 8.
4. Select the status. The interface is enabled by default.
5. Select Add to add the new interface.
6. Select IP.
7. Select IP Address to display the Router IP Address panel.

NOTE: If a table listing the configured IP addresses is displayed instead, click Add IP Address.

8. Select the loopback interface from the Port field's pulldown menu. For example, to select loopback interface 1, select "lb1".
9. Enter the IP address and sub-net mask.
10. Select Add to add the interface.

Adding BGP4 Neighbors

The BGP4 protocol does not contain a peer discovery process. Therefore, for each of the routing switch's BGP4 neighbors (peers), you must indicate the neighbor's IP address and the AS each neighbor is in. Neighbors that are in different ASs communicate using EBGP. Neighbors within the same AS communicate using IBGP.

By default, the routing switch can have up to three BGP neighbors. If you need more support for more neighbors, you can allocate additional CPU memory for neighbors using the **max-neighbors** command. (See "Changing the Maximum Number of Neighbors" on page 12-17.)

USING THE CLI

To add a BGP4 neighbor with IP address 209.157.22.26, enter the following command:

```
HP9300(config-bgp-router)# neighbor 209.157.22.26
```

The neighbor's <IP addr> must be a valid IP address.

The **neighbor** command has some additional parameters, as shown in the following syntax:

Syntax: neighbor <router ID> [remote-as <AS number>]
[advertisement-interval <num>] [distribute-list in|out <num,num,...>]
[ebgp-multihop] [filter-list in|out <num,num,...>] [maximum-prefix <num>]
[next-hop-self] [remote-as <AS number>] [remove-private-as]
[route-map <in|out> <map name>] [route-reflector-client] [send-community]
[update-source loopback <num>] [weight <num>]

advertisement-interval <num> specifies the minimum delay (in seconds) between messages to the specified neighbor. The default is 30 for EBGP neighbors (neighbors in other ASs). The default is 5 for IBGP neighbors (neighbors in the same AS). The range is 0 – 600.

NOTE: The routing switch applies the advertisement interval only under certain conditions. The routing switch does not apply the advertisement interval when sending initial updates to a BGP4 neighbor. As a result, when a routing switch needs to send its entire routing table to a BGP4 neighbor, the routing switch sends the updates one immediately after another at rate of one TCP window per second, without waiting for the advertisement interval.

The routing switch still applies the advertisement interval to an update if the update contains a route for which the routing switch has just sent an update. For example, if the routing switch sends an update for routes 1,2, and 3, then receives a change to an attribute of one of the routes before the advertisement interval has expired, the routing switch waits to send an update for the change until the advertisement interval has expired.

distribute-list in|out <num,num,...> specifies a distribute list to be applied to updates to or from the specified neighbor. The **in|out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor. The <num,num,...> parameter specifies the list of address-list filters. The routing switch applies the filters in the order in which you list them and stops applying the filters in the distribute list when a match is found.

NOTE: By default, if a route does not match any of the filters, the routing switch denies the route. To change the default behavior, configure the last filter as "permit any any".

NOTE: The address filter must already be configured. See “Filtering Specific IP Addresses” on page 12-30.

ebgp-multihop specifies that the neighbor is more than one hop away and that the session type with the neighbor is thus EBGP-multihop. This option is disabled by default.

filter-list in|out <num,num,...> specifies an AS-path filter list. The **in|out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor. The <num,num,...> parameter specifies the list of AS-path filters. The routing switch applies the filters in the order in which you list them and stops applying the filters in the AS-path filter list when a match is found.

NOTE: By default, if an AS-path does not match any of the filters, the routing switch denies the route. To change the default behavior, configure the last filter as "permit any any".

NOTE: The AS-path filter must already be configured. See “Filtering AS-Paths” on page 12-32.

maximum-prefix <num> specifies the maximum number of IP network prefixes (routes) that can be learned from the specified neighbor. The default is 80000 for chassis devices using Redundant Management modules. The default is 5000 for the 6308M-SX and chassis devices using other management modules. The range is from 100 to the maximum number of BGP4 routes allowed on the routing switch. The maximum value depends on the type of routing switch you have and also on whether you have changed the maximum number of routes for the device. See “Changing the Maximum Number of Routes” on page 12-18.

next-hop-self specifies that the routing switch should list itself as the next hop in updates sent to the specified neighbor. This option is disabled by default.

remote-as <AS number> specifies the AS the remote neighbor is in. The <AS number> can be a number from 1 – 65535. There is no default.

remove-private-as configures the routing switch to remove private AS numbers from UPDATE messages the routing switch sends to this neighbor. The routing switch will remove AS numbers 64512 – 65535 (the well-known BGP4 private AS numbers) from the AS-path attribute in UPDATE messages the routing switch sends to the neighbor. This option is disabled by default.

route-map in|out <map name> specifies a route map the routing switch will apply to updates sent to or received from the specified neighbor. The **in|out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor.

NOTE: The route map must already be configured. See “Defining Route Maps” on page 12-36.

route-reflector-client specifies that this neighbor is a route-reflector client of the routing switch. Use the parameter only if this routing switch is going to be a route reflector. For information, see "Configuring Route Reflection Parameters" on page 12-24. This option is disabled by default.

send-community enables sending the community attribute in updates to the specified neighbor. By default, the routing switch does not send the community attribute. This option is disabled by default.

update-source loopback <num> configures the routing switch to communicate with the neighbor through the loopback address on the specified interface. Using a loopback address for neighbor communication avoids problems that can be caused by unstable routing switch interfaces. Generally, loopback interfaces are used for links to IBGP neighbors, which often are multiple hops away, rather than EBGP neighbors. The <num> parameter indicates the loopback interface number and can be from 1 – 4. There is no default.

weight <num> specifies a weight the routing switch will add to routes received from the specified neighbor. BGP4 prefers larger weights over smaller weights. The default weight is 0.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Neighbor link from the bottom of the BGP configuration sheet to display the BGP Neighbor panel, as shown in Figure 12.5.

   NOTE: If you have already defined a neighbor, a table such as the one shown in Figure 12.17 is displayed instead. In this case, select the Add BGP Neighbor link from the bottom of the panel.

2. Enter the neighbor's IP address in the IP Address field.
3. Optionally enable default originate. By default, the routing switch does not advertise a default route using BGP4. A BGP4 default route is the IP address 0.0.0.0 and the route prefix 0 or network mask 0.0.0.0. For example, 0.0.0.0/0 is a default route.
4. If the neighbor is in a remote AS (an EBGP neighbor) and is multiple hops away from the routing switch, enable EBGP Multihop.
5. If you want the routing switch to list itself as the next hop in updates sent to the neighbor, select Next Hop Self.
6. If you want the routing switch to include the community attribute in updates to the neighbor, enable Send Community.
7. If you want the routing switch to remove the private AS from updates to the neighbor, enable Remove Private AS.
8. Optionally change the advertisement interval by changing the value in the Advert Interval field. The default is 30 for EBGP neighbors (neighbors in other ASs). The default is 5 for IBGP neighbors (neighbors in the same AS). The range is 0 – 600.
9. Optionally change the maximum number of IP network prefixes (routes) that can be learned from the neighbor. The default is 5000. The range is from 100 to the maximum number of routes allowed on the routing switch. The maximum value depends on the type of routing switch you have and also on whether you have changed the maximum number of routes for the device. See "Changing the Maximum Number of Routes" on page 12-18.
10. Enter the AS number for the neighbor in the Remote AS field. (This field also applies to local neighbors. For a local neighbor, enter the local AS.)
11. Optionally add a weight to routes received from the neighbor. BGP4 prefers routes with larger weights over routes with smaller weights. The default weight is 0.
12. Optionally specify a loopback interface number in the update source in the Update Source field. Using this option configures the routing switch to communicate with the neighbor through the loopback address on the specified interface. Using a loopback address for neighbor communication avoids problems that can be caused by unstable routing switch interfaces. Generally, loopback interfaces are used for links to IBGP neighbors, which often are multiple hops away, rather than EBGP neighbors. The loopback interface number can be from 1 – 4. There is no default.
13. Optionally enter a list of AS-path filters.

14. Click the Add button to assign the changes.

NOTE: You need to configure the AS-path filters before you can apply them to the neighbor. See “Filtering AS-Paths” on page 12-32.

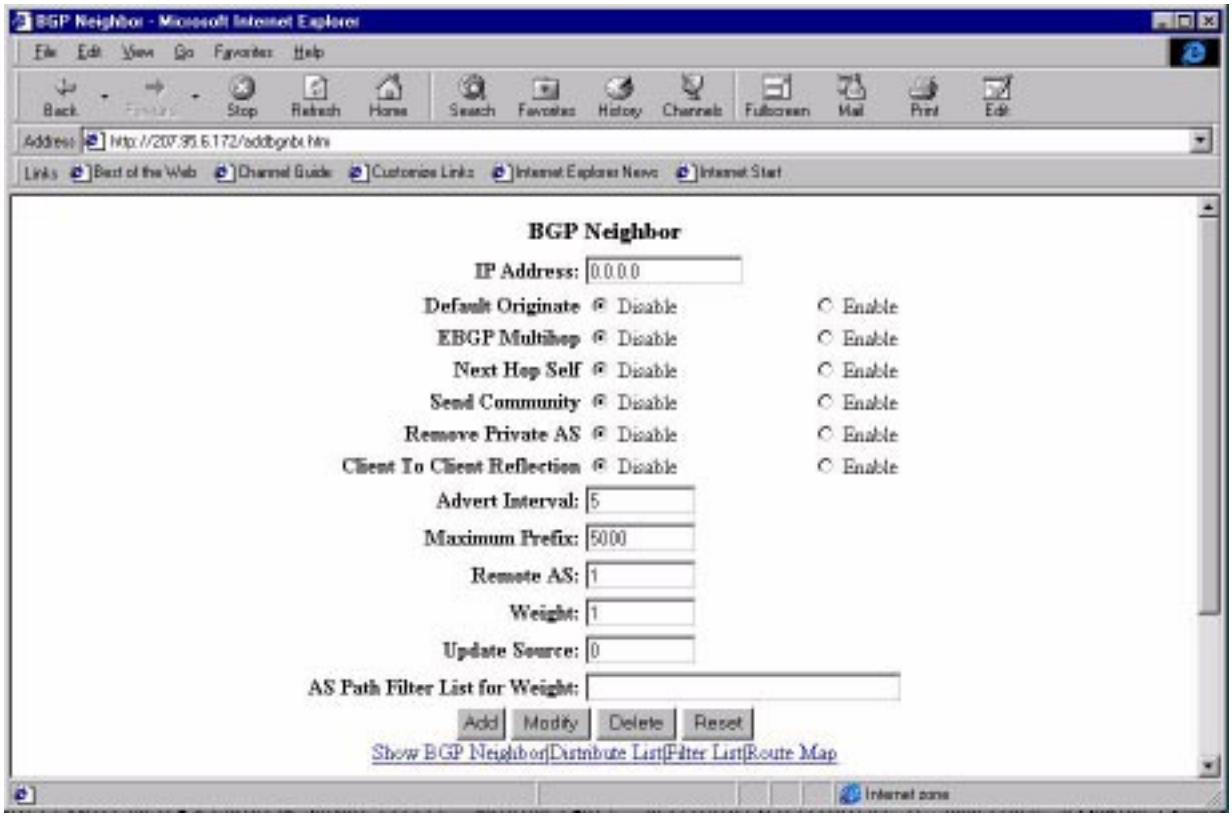

Figure 12.5 BGP4 Neighbor panel

Optional Configuration Tasks

The following sections describe how to perform optional BGP4 configuration tasks.

Changing the Keep Alive Time and Hold Time

The Keep Alive Time specifies how frequently the routing switch will send KEEPALIVE messages to its BGP4 neighbors. The Hold Time specifies how long the routing switch will wait for a KEEPALIVE or UPDATE message from a neighbor before concluding that the neighbor is dead. When the routing switch concludes that a BGP4 neighbor is dead, the routing switch ends the BGP4 session and closes the TCP connection to the neighbor.

The default Keep Alive time for routing switches is 60 seconds. The default Hold Time is 180 seconds. To change the timers, use either of the following methods.

NOTE: Generally, you should set the Hold Time to three times the value of the Keep Alive Time.

USING THE CLI

To change the Keep Alive Time to 30 and Hold Time to 90, enter the following command:

```
HP9300(config-bgp-router)# timers keep-alive 30 hold-time 90
```

Syntax: timers keep-alive <num> hold-time <num>

For each keyword, <num> indicates the number of seconds. The Keep Alive Time can be 0 – 65535. The Hold Time can be 0 or 3 – 65535 (1 and 2 are not allowed). If you set the Hold Time to 0, the routing switch waits indefinitely for messages from a neighbor without concluding that the neighbor is dead.

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Edit the number in the Keep Alive Time field.
3. Edit the number in the Hold Time field.
4. Select Apply to implement the change.

Disabling Fast External Fallover

BGP4 routers rely on KEEPALIVE and UPDATE messages from neighbors to signify that the neighbors are alive. For BGP4 neighbors that are two or more hops away, such messages are the only indication that the BGP4 protocol has concerning the alive state of the neighbors. As a result, if a neighbor dies, the routing switch will wait until the Hold Time expires before concluding that the neighbor is dead and closing its BGP4 session and TCP connection with the neighbor.

For directly attached neighbors, the routing switch immediately senses loss of a connection to the neighbor from a change to the state of the port or interface that connects the routing switch to its neighbor. For directly attached EBGP neighbors, the routing switch can use this information to immediately close the BGP4 session and TCP connection to locally attached neighbors that die.

NOTE: The fast external fallover feature applies only to directly attached EBGP neighbors. The feature does not apply to IBGP neighbors.

If you want to disable this behavior and thus configure the routing switch to wait for the Hold Time to expire before ending the connection to a directly-attached BGP4 neighbor that dies, use either of the following methods.

USING THE CLI

To disable fast external fallover, enter the following command:

```
HP9300(config-bgp-router)# no fast-external-fallover
```

syntax: [no] fast-external-fallover

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Select Disable next to Fast External Fall Over.
3. Select Apply to implement the change.

Changing the Maximum Number of Neighbors

The 9304M, 9308M, and 6308M-SX routing switches can have up to three BGP4 neighbors by default. You can change the maximum number of BGP4 neighbors the routing switch can have from 1 – 10 using either of the following methods.

NOTE: If you have a lot of IBGP neighbors, you can configure some IBGP routers as route reflectors. By doing so, you can reduce the number of neighbors you need to configure on each routing switch. Without route reflectors, all IBGP routers must be fully meshed to ensure proper route propagation. See “Configuring Route Reflection Parameters” on page 12-24.

USING THE CLI

To change the maximum number of BGP4 neighbors to 10, enter the following command:

```
HP9300(config-bgp-router)# max-neighbors 10
HP9300(config-bgp-router)# exit
HP9300# reload
```

Syntax: max-neighbors <num>

The <num> indicates the number of BGP4 neighbors allowed and can be from 1 – 10. The change takes effect after the routing switch is rebooted.

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Change the number in the Maximum Neighbor field. You can enter a number from 1 – 10.
3. Select Apply to implement the change.
4. Select the Reload link to reload the software. You must reload or reboot to place the change into effect.

Changing the Maximum Number of Routes

The BGP4 route table holds 30,000 routes by default and can hold up to 70,000 routes. You can change the maximum number of BGP4 routes the routing switch can have using either of the following methods.

NOTE: On 9304M and 9308M routing switches using Redundant Management modules, the default maximum number of routes is 256,000.

NOTE: This value also determines the maximum value you can configure when specifying how many routes this routing switch can learn from a neighbor. See the description of the maximum prefix option in “Adding BGP4 Neighbors” on page 12-13.

USING THE CLI

To change the maximum number of BGP4 routes to 60000, enter the following command:

```
HP9300(config-bgp-router)# max-routes 60000
HP9300(config-bgp-router)# exit
HP9300# reload
```

Syntax: max-routes <num>

The <num> indicates the number of BGP4 routes allowed and can be from 10000 – 70000. The change takes effect after the routing switch is rebooted.

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Change the number in the Maximum Route field. You can enter a number from 10000 – 70000.
3. Select Apply to implement the change.
4. Select the Reload link to reload the software. You must reload or reboot to place the change into effect.

Changing the Maximum Number of Route-Attribute Entries

The BGP4 route table lists the route attributes associated with each route in the table. These attributes include the following:

- IP address of the next hop router
- Metric

- Local Preference
- Weight
- Origin
- Static
- Route tag
- Communities

A collection of these attributes is called a ***route-attributes entry***. Each route-attributes entry is a unique set of values for these attributes. For example, the following set of attribute values is a route-attributes entry:

```
Next Hop  :192.168.11.1       Metric   :0               Origin:IGP
Originator:0.0.0.0            Cluster List:None
Aggregator:AS Number :0       Router-ID:0.0.0.0        Atomic:FALSE
Local Pref:100                Communities:Internet
```

A route-attribute entry can be used by one or more routes. For example, if the first and second routes listed in the BGP4 route table use exactly the same set of attribute values, the routes both would use a single route-attributes entry. If any of the attributes differs for the two routes, each route would use a separate route-attributes entry. See “Displaying BGP4 Route-Attribute Entries” on page 12-55 for a description of the route-attribute fields shown in the example above.

By default, the routing switch can contain 10000 route-attribute entries. You can change the maximum number of route-attribute entries the routing switch can contain from 200 – 30000 using either of the following methods.

NOTE: On 9304M and 9308M routing switches using Redundant Management modules, the default maximum number of route-attribute entries is 100,000.

USING THE CLI

To change the maximum number of route-attribute entries to 25000, enter the following command:

```
HP9300(config-bgp-router)# max-attribute-entries 25000
HP9300(config-bgp-router)# exit
HP9300# reload
```

Syntax: max-attribute-entries <num>

The <num> indicates the number of route-attribute entries allowed on the routing switch. The value can be from 200 – 30000. The change takes effect after the routing switch is rebooted.

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Change the number in the Maximum Attribute Entries field. You can enter a number from 200 – 30000.
3. Select Apply to implement the change.
4. Select the Reload link to reload the software. You must reload or reboot to place the change into effect.

Specifying a List of Networks to Advertise

By default, the routing switch sends BGP4 routes only for the networks you identify using the **network** command or that are redistributed into BGP4 from RIP or OSPF. To specify a network to be advertised, use either of the following methods.

USING THE CLI

To configure the routing switch to advertise network 209.157.22.0/24, enter the following command:

```
HP9300(config-bgp-router)# network 209.157.22.0 255.255.255.0
```

syntax: network <IP-addr> mask <network-mask> [weight <num>] [backdoor]

The <IP-addr> is the network number and the **mask** <network-mask> specifies the network mask.

The **weight** <num> parameter specifies a weight to be added to routes to this network.

The **backdoor** parameter changes the administrative distance of the route to this network from the EBGP administrative distance (20 by default) to the Local BGP weight (200 by default), thus tagging the route as a backdoor route. Use this parameter when you want the routing switch to prefer IGP routes such as RIP or OSPF routes over the EBGP route for the network.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Network link from the bottom of the BGP configuration sheet to display the BGP Network panel, shown in Figure 12.6.
2. Enter the network address in the IP Address field.
3. Enter the network mask in the Mask field.
4. Optionally enter a weight to be added to routes to this network.
5. If you want to tag the route as a backdoor route, select Enable next to Back Door.
6. Select Apply to implement the change.

The backdoor parameter changes the administrative distance of the route to this network from the EBGP administrative distance (20 by default) to the Local BGP weight (200 by default), thus tagging the route as a backdoor route. Use this parameter when you want the routing switch to prefer IGP routes such as RIP or OSPF routes over the EBGP route for the network.

Figure 12.6 Network panel

Changing the Default Local Preference

When the routing switch uses the BGP4 algorithm to select a route to send to the IP route table, one of the parameters the algorithm uses is the local preference. Local preference is an attribute that indicates a degree of preference for a route relative to other routes in the local AS. BGP4 neighbors can send the local preference value as an attribute of a route in an UPDATE message.

Local preference applies only to routes within the local AS. BGP4 routers can exchange local preference information with neighbors who also are in the local AS, but BGP4 routers do not exchange local preference information with neighbors in remote ASs.

The default local preference is 100. When the BGP4 algorithm compares routes on the basis of local preferences, the route with the higher local preference is chosen.

NOTE: To set the local preference for individual routes, use route maps. See “Defining Route Maps” on page 12-36. See “How BGP4 Selects a Path for a Route” on page 12-3 for information about the BGP4 algorithm.

To change the default local preference used by the routing switch, use either of the following methods.

USING THE CLI

To change the default local preference to 200, enter the following command:

```
HP9300(config-bgp-router)# default-local-preference 200
```

Syntax: default-local-preference <num>

The <num> parameter indicates the preference and can be a value from 0 – 4294967295.

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Change the number in the Default Local Preference field. You can enter a number from 0 – 4294967295.
3. Select Apply to implement the change.

Advertising the Default Information Originate

By default, the routing switch does not advertise a default route using BGP4. A BGP4 default route is the IP address 0.0.0.0 and the route prefix 0 or network mask 0.0.0.0. For example, 0.0.0.0/0 is a default route. You can enable the routing switch to advertise a default BGP4 route using either of the following methods.

USING THE CLI

To enable the routing switch to advertise a default BGP4 route, enter the following command:

```
HP9300(config-bgp-router)# default-information-originate
```

Syntax: [no] default-information-originate

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Select Enable next to Default Information Originate.
3. Select Apply to implement the change.

Changing the Default MED (Metric) Used for Route Redistribution

The routing switch can redistribute RIP and OSPF routes into BGP4. The MED (metric) is a global parameter that specifies the cost that will be applied to all routes by default when they are redistributed into BGP4. When routes are selected, lower metric values are preferred over higher metric values. The default BGP4 MED value is 0 and can be assigned a value from 0 – 4294967295.

NOTE: RIP and OSPF also have default metric parameters. The parameters are set independently for each protocol and have different ranges.

USING THE CLI

To change the default metric to 40, enter the following command:

```
HP9300(config-bgp-router)# default-metric 40
```

Syntax: default-metric <num>

The <num> indicates the metric and can be a value from 0 – 4294967295.

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Change the number in the Default Metric field. You can enter a number from 0 – 4294967295.
3. Select Apply to implement the change.

Changing Administrative Distances

BGP4 routers can learn about networks from various protocols, including the EBGP portion of BGP4 and IGPs such as OSPF and RIP. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned.

To select one route over another based on the source of the route information, the routing switch can use the administrative distances assigned to the sources.

NOTE: The software will replace a statically configured static default route with a learned default route if the learned route's administrative distance is lower than the statically configured default route's distance. However, the default administrative distance for static routes is 1, so only directly-connected routes are preferred over static routes when the default administrative distances for the routes are used.

Here are the default administrative distances on the 9304M, 9308M, and 6308M-SX routing switches:

- Directly connected – 0 (this value is not configurable)
- Static – 1 (applies to all static routes, including default routes)
- EBGP – 20
- OSPF – 110
- RIP – 120
- IBGP – 200
- Local BGP – 200
- Unknown – 255 (the routing switch will not use this route)

Lower administrative distances are preferred over higher distances. For example, if the routing switch receives routes for the same network from OSPF and from RIP, the routing switch will prefer the OSPF route by default. The administrative distances are configured in different places in the software.

- To change the EBGP, IBGP, and Local BGP default administrative distances, see the instructions in this section.
- To change the default administrative distance for OSPF, see "Modify Administrative Distance" on page 10-26.
- To change the default administrative distance for RIP, see "Modifying the Default Administrative Distance" on page 9-31.
- To change the default administrative distance for static routes, see "Defining Static IP Routes" on page 9-11.

You can change the default EBGP, IBGP, and Local BGP administrative distances using either of the following methods.

USING THE CLI

To change the default administrative distances for EBGP, IBGP, and Local BGP, enter a command such as the following:

```
HP9300(config-bgp-router)# distance 180 160 40
```

syntax: distance <external-distance> <internal-distance> <local-distance>

The <external-distance> sets the EBGP distance and can be a value from 1 – 255.

The <internal-distance> sets the IBGP distance and can be a value from 1 – 255.

The <local-distance> sets the Local BGP distance and can be a value from 1 – 255.

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Change the number in the External Distance field to change the EBGP distance. You can enter a number from 1 – 255.

3. Change the number in the Internal Distance field to change the IBGP distance. You can enter a number from 1 – 255.
4. Change the number in the Local Distance field to change the local distance. You can enter a number from 1 – 255.
5. Select Apply to implement the change.

Configuring the routing switch To Always Compare Multi-Exit Discriminators (MEDs)

A Multi-Exit Discriminator (MED) is a value that the BGP4 algorithm uses when comparing multiple paths received from different BGP4 neighbors in the same AS for the same route. In BGP4, a route's MED is equivalent to its "metric".

By default, the routing switch compares the MED values only among paths through the same AS. For example, if the routing switch receives BGP4 updates from a remote AS with multiple paths for the same route, the routing switch compares the MEDs in those paths to select a preferred path for the route.

You can change the routing switch's default behavior and configure the routing switch to instead compare the MEDs for all paths for a route, regardless of the AS through which the paths pass. For example, if the routing switch receives UPDATES for the same route from neighbors in three ASs, the routing switch would compare the MEDs of all the paths together, rather than comparing the MEDs for the paths in each AS individually.

To configure the routing switch to always compare MEDs for all paths for a route, use either of the following methods:

USING THE CLI

To configure the routing switch to always compare MEDs, enter the following command:

```
HP9300(config-bgp-router)# always-compare-med
```

syntax: [no] always-compare-med

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Select Enable next to Always Compare MED.
3. Select Apply to implement the change.

Synchronizing Routes

By default, the routing switch waits until the IGPs in the local AS have fully exchanged route information before BGP4 advertises the routes to its remote BGP4 neighbors. This behavior is called synchronization.

If you instead want the routing switch to advertise routes to its remote BGP4 neighbors regardless of whether the routes are learned or have already been propagated throughout the local AS, you can disable synchronization. To disable synchronization, use either of the following methods.

USING THE CLI

To disable synchronization, enter the following command:

```
HP9300(config-bgp-router)# no synchronization
```

To re-enable synchronization, enter the following command:

```
HP9300(config-bgp-router)# synchronization
```

Syntax: [no] synchronization

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Select Disable or Enable next to Synchronization.
3. Select Apply to implement the change.

Automatically Summarizing Subnet Routes Into Class A, B, or C Networks

The auto summary feature summarizes the routes it redistributes from IGP to BGP4. The routing switch summarizes subnets into their natural class A, B, or C networks. For example, if an AS contains subnets 1.1.0.0, 1.2.0.0, and 1.3.0.0 with the network mask 255.255.0.0, the auto summary feature summarizes the subnets in its advertisements to BGP4 neighbors as 1.0.0.0/8.

The auto summary feature is enabled by default. If you want to disable the feature, use either of the following methods.

NOTE: The auto summary feature summarizes only the routes that are redistributed from IGP into BGP4.

NOTE: The auto summary feature does not summarize networks that use CIDR numbers instead of class A, B, or C numbers. To summarize CIDR networks, use the aggregation feature. See “Aggregating Routes Advertised to BGP4 Neighbors” on page 12-26.

USING THE CLI

To disable auto summary, enter the following command:

```
HP9300(config-bgp-router)# no auto-summary
```

To re-enable auto summary, enter the following command:

```
HP9300(config-bgp-router)# auto-summary
```

syntax: [no] auto-summary

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP configuration sheet, shown in Figure 12.3.
2. Select Disable or Enable next to Auto Summary.
3. Select Apply to implement the change.

Configuring Route Reflection Parameters

Normally, all the BGP routers within an AS are fully meshed. Each of the routers has an IBGP session with each of the other BGP routers in the AS. Each IBGP router thus has a route for each of its IBGP neighbors. For large ASs containing 100 or more IBGP routers, the IBGP route information in each of the fully-meshed IBGP routers can introduce too much administrative overhead.

To avoid this problem, you can hierarchically organize your IGP routers into clusters.

- A ***cluster*** is a group of IGP routers organized into route reflectors and route reflector clients. You configure the cluster by assigning a cluster ID on the route reflector and identifying the IGP neighbors that are members of that cluster. All the configuration for route reflection takes place on the route reflectors. The clients are unaware that they are members of a route reflection cluster. All members of the cluster must be in the same AS. The cluster ID can be any number from 1 – 4294967295. The default is the router ID, expressed as a 32-bit number.

  NOTE: If the cluster contains more than one route reflector, you need to configure the same cluster ID on all the route reflectors in the cluster. The cluster ID helps route reflectors avoid loops within the cluster.

- A ***route reflector*** is an IGP router configured to send BGP route information to all the clients (other BGP4 routers) within the cluster. Route reflection is enabled on the 9304M, 9308M, and 6308M-SX routing switches by default but does not take effect unless you add route reflector clients and cluster ID information to the routing switch.
- A ***route reflector client*** is an IGP router identified as a member of a cluster. You identify a router as a route reflector client on the routing switch that is the route reflector, not on the client. The client itself requires no additional configuration. In fact, the client does not know that it is a route reflector client. The client just knows that it receives updates from its neighbors and does not know whether one or more of those neighbors are route reflectors.

NOTE: Route reflection applies only among IBGP routers within the same AS. You cannot configure a cluster that spans multiple ASs.

Figure 12.7 shows an example of a route reflector configuration. In this example, two routing switches are configured as route reflectors for the same cluster. The route reflectors provide redundancy in case one of the reflectors becomes unavailable. Without redundancy, if a route reflector becomes unavailable, its clients are cut off from BGP4 updates.

AS1 contains a cluster with two route reflectors and two clients. The route reflectors are fully meshed with other BGP4 routers, but the clients are not fully meshed. They rely on the route reflectors to propagate BGP4 route updates.

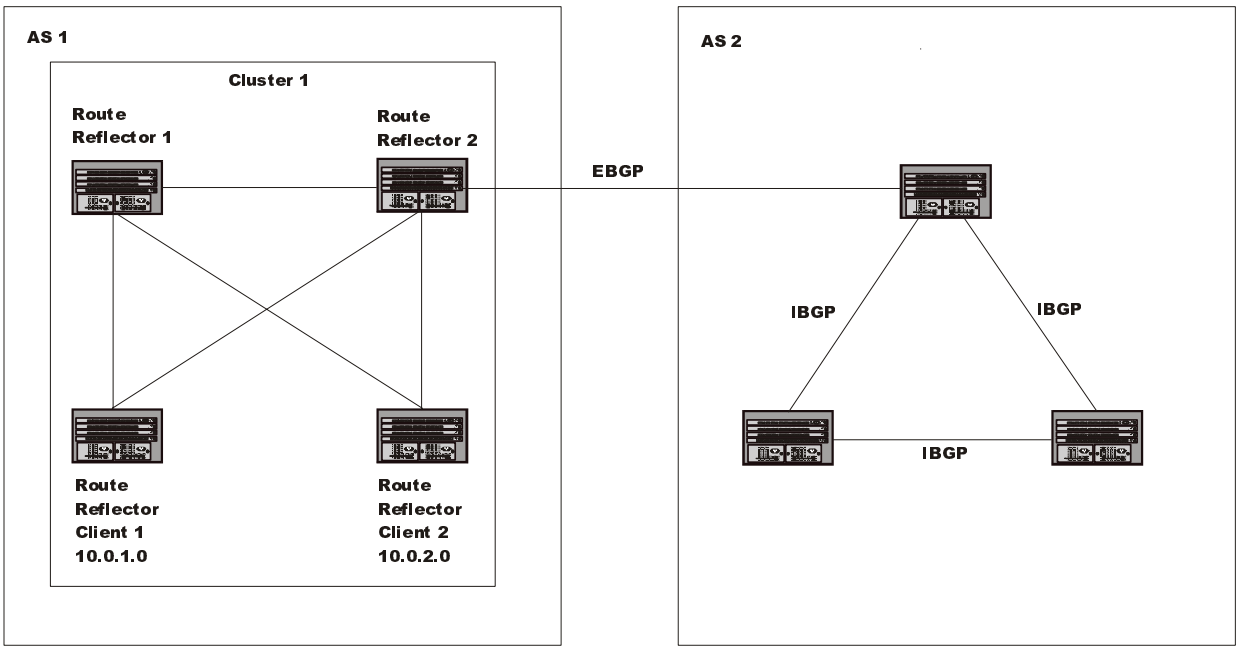

Figure 12.7 Example route reflector configuration

To configure a 9304M, 9308M, or 6308M-SX routing switch to be a BGP4 route reflector, use either of the following methods.

NOTE: All configuration for route reflection takes place on the route reflectors, not on the clients.

USING THE CLI

Enter the following commands to configure a routing switch as route reflector 1 in Figure 12.7. To configure route reflector 2, enter the same commands on the routing switch that will be route reflector 2. The clients require no configuration for route reflection.

```
HP9300(config-bgp-router)# cluster-id 1
HP9300(config-bgp-router)# neighbor 10.0.1.0 route-reflector-client
HP9300(config-bgp-router)# neighbor 10.0.2.0 route-reflector-client
```

syntax: cluster-id <num>

The <num> parameter specifies the cluster ID and can be a number from 1 – 4294967295. The default is the router ID, expressed as a 32-bit number. You can configure one cluster ID on the routing switch. All route-reflector clients for the routing switch are members of the cluster.

NOTE: If the cluster contains more than one route reflector, you need to configure the same cluster ID on all the route reflectors in the cluster. The cluster ID helps route reflectors avoid loops within the cluster.

To add an IBGP neighbor to the cluster, enter the following command:

syntax: neighbor <IP-addr> route-reflector-client

For more information abut the **neighbor** command, see “Adding BGP4 Neighbors” on page 12-13.

If you need to disable route reflection on a routing switch, enter the following command. Disabling route reflection allows you to turn off the feature without removing Cluster ID and route reflector client information from the system configuration file.

```
HP9300(config-bgp-router)# no client-to-client-reflection
```

Enter the following command to re-enable the feature:

```
HP9300(config-bgp-router)# client-to-client-reflection
```

syntax: [no] client-to-client-reflection

USING THE WEB MANAGEMENT INTERFACE

1. Select the BGP link to display the BGP panel, shown in Figure 12.5.
2. Client-to-client reflection (route reflection) is enabled by default. If the feature has been disabled, re-enable it by selecting Enable next to Client To Client Reflection.
3. If the autonomous system (AS) the routing switch is in will contain more than one route reflector (a route reflector in addition to the routing switch), enter a cluster ID in the Cluster ID field. The cluster ID is required to avoid loops in an AS that contains more than one route reflector.
4. Click the Apply button to apply the changes.
5. Select the Neighbor link from the list of BGP4 links at the bottom of the panel. This link displays the Neighbor panel.
6. If you have already configured neighbors, a table listing the neighbors is displayed. Click Modify next to the neighbor you want to identify as a route reflector client or select the Add Neighbor link to display the panel shown below.

   If you have not already configured a neighbor, the panel shown below is displayed.
7. Enter the IP address of the neighbor.
8. Configure or change other parameters if needed, then identify this neighbor as a route reflector client by selecting Enable next to Client To Client Reflection.
9. Click the Add button to apply the change.
10. Select the Save To Flash link, then select Yes when prompted to save the changes.

Aggregating Routes Advertised to BGP4 Neighbors

By default, the routing switch advertises individual routes for all the networks. The aggregation feature allows you to configure the routing switch to aggregate routes in a range of networks into a single CIDR number. For example, without aggregation, the routing switch will individually advertise routes for networks 207.95.1.0, 207.95.2.0, and 207.95.3.0. You can configure the routing switch to instead send a single, aggregate route for the networks. The aggregate route would be advertised as 207.95.0.0.

NOTE: To summarize CIDR networks, you must use the aggregation feature. The auto summary feature does not summarize networks that use CIDR numbers instead of class A, B, or C numbers.

To aggregate routes, use either of the following methods.

USING THE CLI

To aggregate routes for 209.157.22.0, 209.157.23.0, and 209.157.24.0, enter the following command:

```
HP9300(config-bgp-router)# aggregate-address 209.157.0.0 255.255.0.0
```

syntax: aggregate-address <IP-addr> <network-mask> [as-set] [summary-only] [suppress-map <map-name>] [advertise-map <map-name>] [attribute-map <map-name>]

The <IP-addr> and <network-mask> parameters specify the aggregate value for the networks. Specify 0 for the host portion and for the network portion that differs among the networks in the aggregate. For example, to aggregate 10.0.1.0, 10.0.2.0, and 10.0.3.0, enter the IP address 10.0.0.0 and the network mask 255.255.0.0.

The **as-set** parameter causes the routing switch to aggregate AS-path information for all the routes in the aggregate address into a single AS-path.

The **summary-only** parameter prevents the routing switch from advertising more specific routes contained within the aggregate route.

The **suppress-map** <map-name> parameter prevents the more specific routes contained in the specified route map from being advertised.

The **advertise-map** <map-name> parameter configures the routing switch to advertise the more specific routes in the specified route map.

The **attribute-map** <map-name> parameter configures the routing switch to set attributes for the aggregate routes based on the specified route map.

NOTE: For the **suppress-map**, **advertise-map**, and **attribute-map** parameters, the route map must already be defined. See “Defining Route Maps” on page 12-36 for information on defining a route map.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Aggregate Address link at the bottom of the BGP configuration sheet to display the panel shown in Figure 12.3.
2. Enter the aggregate address in the IP Address field. Specify 0 for the host portion and for the network portion that differs among the networks in the aggregate. For example, to aggregate 10.0.1.0, 10.0.2.0, and 10.0.3.0, enter the IP address 10.0.0.0. Then enter 255.255.0.0 in the Mask field.
3. Enter the mask in the Mask field.
4. Select one of the following options from the Option field’s pulldown list:
   - Address – Use this option when you are adding the address. This is the default option.
   - AS Set – This option causes the routing switch to aggregate AS-path information for all the routes in the aggregate address into a single AS-path.
   - Summary Only – This option prevents the routing switch from advertising more specific routes contained within the aggregate route.
   - Suppress Map – This option prevents the more specific routes contained in the specified route map from being advertised.
   - Advertise Map – This option configures the routing switch to advertise the more specific routes in the specified route map.
   - Attribute Map – This option configures the routing switch to set attributes for the aggregate routes based on the specified route map.
5. Optionally select a route map from the Map field’s pulldown list.

NOTE: For the Suppress Map, Advertise Map, and Attribute Map options, you must select a route map and the route map must already be defined. See “Defining Route Maps” on page 12-36 for information on defining a route map.

6. Select Add to implement the change.

Figure 12.8 BGP Aggregate Address panel

Modifying Redistribution Parameters

By default, the routing switch does not redistribute route information between BGP4 and the IP IGPs (RIP and OSPF). You can configure the routing switch to redistribute OSPF routes, RIP routes, or static routes into BGP4. The following subsections describe how to set redistribution parameters.

Redistributing RIP Routes

USING THE CLI

To configure BGP4 to redistribute RIP routes and add a weight of 10 to the redistributed routes, enter the following command:

```
HP9300(config-bgp-router)# redistribute rip weight 10
```

syntax: redistribute rip [metric <num>] [route-map <map-name>] [weight <num>]

The **rip** parameter indicates that you are redistributing RIP routes into BGP4.

The **metric** <num> parameter changes the metric. You can specify a value from 0 – 4294967295. The default is 0.

The **route-map** <map-name> parameter specifies a route map to be consulted before adding the filter to the IP route table.

NOTE: The route map you specify must already be configured on the routing switch. See “Defining Route Maps” on page 12-36 for information about defining route maps.

The **weight** <num> parameter changes the weight. You can specify a value from 0 – 65535. The default is 0.

USING THE WEB MANAGEMENT INTERFACE

The following procedure applies to redistributing RIP, OSPF, and Static routes.

1. Select the Redistribute link from the bottom of the BGP configuration sheet to display the BGP Redistribute panel, shown in Figure 12.9.
2. Select the source of the routes you want to redistribute into BGP4. You can select RIP, OSPF, or Static.
3. Optionally enter a metric for the redistributed routes in the Metric field. You can specify a value from 0 – 4294967295. The default is 0.
4. Optionally select a route map from the Map field's pulldown list.

   NOTE: The route map must already be defined. See "Defining Route Maps" on page 12-36 for information on defining a route map.

5. Optionally enter a weight for the redistributed routes in the Weight field. You can specify a value from 0 – 65535. The default is 0.
6. For OSPF routes, select one of the following to specify the types of OSPF routes to be redistributed into BGP4:
   - Internal
   - External
   - External 2
7. Select Add to implement the change.

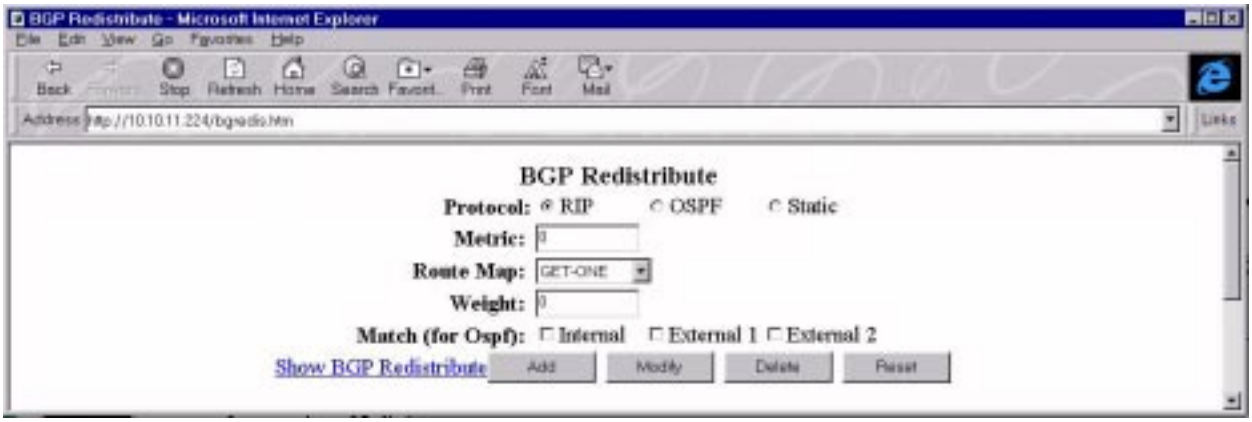

Figure 12.9 BGP Redistribute panel

Redistributing OSPF Routes

To configure the routing switch to redistribute OSPF external type 1 routes, enter the following command:

```
HP9300(config-bgp-router)# redistribute ospf match external1
```

syntax: redistribute ospf [metric <num>] [route-map <map-name>] [weight <num>] [match internal|external1|external2]

The **ospf** parameter indicates that you are redistributing OSPF routes into BGP4.

The **metric** <num> parameter changes the metric. You can specify a value from 0 – 4294967295. The default is 0.

The **route-map** <map-name> parameter specifies a route map to be consulted before adding the filter to the IP route table.

NOTE: The route map you specify must already be configured on the routing switch. See “Defining Route Maps” on page 12-36 for information about defining route maps.

The **weight** <num> parameter changes the weight. You can specify a value from 0 – 65535. The default is 0.

The **match internal|external1|external2** parameter applies only to OSPF. This parameter specifies the types of OSPF routes to be redistributed into BGP4.

USING THE Web Management Interface

Use the procedure in “Redistributing RIP Routes” on page 12-28.

Redistributing Static Routes

To configure the routing switch to redistribute static routes, enter the following command:

```
HP9300(config-bgp-router)# redistribute static
```

syntax: redistribute static [metric <num>] [route-map <map-name>] [weight <num>]

The **static** parameter indicates that you are redistributing static routes into BGP4.

The **metric** <num> parameter changes the metric. You can specify a value from 0 – 4294967295. The default is 0.

The **route-map** <map-name> parameter specifies a route map to be consulted before adding the filter to the IP route table.

NOTE: The route map you specify must already be configured on the routing switch. See “Defining Route Maps” on page 12-36 for information about defining route maps.

The **weight** <num> parameter changes the weight. You can specify a value from 0 – 65535. The default is 0.

USING THE WEB MANAGEMENT INTERFACE

Use the procedure in “Redistributing RIP Routes” on page 12-28.

Filtering Specific IP Addresses

You can configure the routing switch to explicitly permit or deny specific IP addresses received in updates from BGP4 neighbors by defining IP address filters. The routing switch permits all IP addresses by default. You can define up to 100 IP address filters for BGP4.

- If you want permit to remain the default behavior, define individual filters to deny specific IP addresses.
- If you want to change the default behavior to deny, define individual filters to permit specific IP addresses.

NOTE: Once you define a filter, the default action for addresses that do not match a filter is "deny". To change the default action to "permit", configure the last filter as "permit any any".

Address filters can be referred to by a BGP neighbor's distribute list number as well as by match statements in a route map.

To define an IP address filter, use either of the following methods.

USING THE CLI

To define an IP address filter to deny routes to 209.157.0.0, enter the following command:

```
HP9300(config-bgp-router)# address-filter 1 deny 209.157.0.0 255.255.0.0
```

syntax: address-filter <num> permit|deny <IP-addr>|any <network-mask>|any

The <num> parameter identifies the filter’s position in the address filter list and can be from 1 – 100. Thus, the address filter list can contain up to 100 filters. The routing switch applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the routing switch stops and does not continue applying filters from the list.

NOTE: If the filter is referred to by a route map's match statement, the filter is applied in the order in which the filter is listed in the match statement.

The **permit|deny** parameter indicates the action the routing switch takes if the filter match is true.

- If you specify **permit**, the routing switch permits the route into the BGP4 table if the filter match is true.
- If you specify **deny**, the routing switch denies the route from entering the BGP4 table if the filter match is true.

The <IP-addr> <network-mask> parameter indicates the IP address you want to filter. If you specify **any any**, all IP routes containing the specified IP addresses are permitted or denied (assuming the IP address is not filtered by a lower-numbered filter with the opposite action).

NOTE: Once you define a filter, the default action for addresses that do not match a filter is "deny". To change the default action to "permit", configure the last filter as "permit any any".

To filter based on network mask only, enter **any** for the IP address and then enter the network mask.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Address Filter link from the bottom of the BGP configuration sheet to display the BGP Address Filter panel, shown in Figure 12.10.
2. Enter the filter ID in the ID field. You can specify a number from 1 – 100.
3. Select the action you want the routing switch to perform if the filter is true.
   - If you select Deny, the routing switch denies the route from entering the BGP4 table if the filter match is true.
   - If you select Permit, the routing switch permits the route into the BGP4 table if the filter match is true.
4. Enter the network prefix in the Prefix field. If you specify "any", all networks match the filter.
5. Enter the prefix masking bits in the Prefix Masking Bits field. The prefix masking bits indicate the bits in the prefix that the filter compares. The filter disregards the bits for which the mask contains zeros.
6. Enter the mask in the Prefix Mask field. If you specify "any", all masks match the filter.
7. Enter the masking bits for the network mask in the Prefix Mask Masking Bits field.
8. Select Add to add the filter.

Figure 12.10 BGP Address Filter panel

Filtering AS-Paths

You can filter updates received from BGP4 neighbors based on the contents of the AS-path list accompanying the updates. For example, if you want to deny routes that have the AS 4.3.2.1 in the AS-path from entering the BGP4 route table, you can define a filter to deny such routes. You can define up to 100 AS-path filters for BGP4.

NOTE: Once you define a filter, the default action for updates that do not match a filter is "deny". To change the default action to "permit", configure the last filter as "permit any any".

AS-path filters can be referred to by a BGP neighbor's distribute list number as well as by match statements in a route map.

To define an AS-path filter, use either of the following methods.

USING THE CLI

To define AS-path filter 4 to permit AS 2500, enter the following command:

```
HP9300(config-bgp-router)# as-path-filter 4 permit 2500
```

syntax: as-path-filter <num> permit|deny <AS-path>

The <num> parameter identifies the filter's position in the AS-path filter list and can be from 1 – 100. Thus, the AS-path filter list can contain up to 100 filters. The routing switch applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the routing switch stops and does not continue applying filters from the list.

NOTE: If the filter is referred to by a route map's match statement, the filter is applied in the order in which the filter is listed in the match statement.

The **permit|deny** parameter indicates the action the routing switch takes if the filter match is true.

- If you specify **permit**, the routing switch permits the route into the BGP4 table if the filter match is true.
- If you specify **deny**, the routing switch denies the route from entering the BGP4 table if the filter match is true.

The <AS-path> parameter indicates the AS-path information. You can enter an exact AS-path string if you want to filter for a specific value. You also can use regular expressions in the filter string.

USING THE WEB MANAGEMENT INTERFACE

1. Select the AS Path Filter link from the bottom of the BGP configuration sheet to display the BGP AS Path Filter panel, shown in Figure 12.11.
2. Enter the filter ID in the ID field. You can specify a number from 1 – 100.
3. Select the action you want the routing switch to perform if the filter is true.
    - If you select Deny, the routing switch denies the route from entering the BGP4 table if the filter match is true.
    - If you select Permit, the routing switch permits the route into the BGP4 table if the filter match is true.
4. Enter the AS path you want to filter in the Regular Expression field. As indicated by the field's title, you can use regular expressions for the AS path. See "Using Regular Expressions" on page 12-33.
5. Select Add to add the filter.

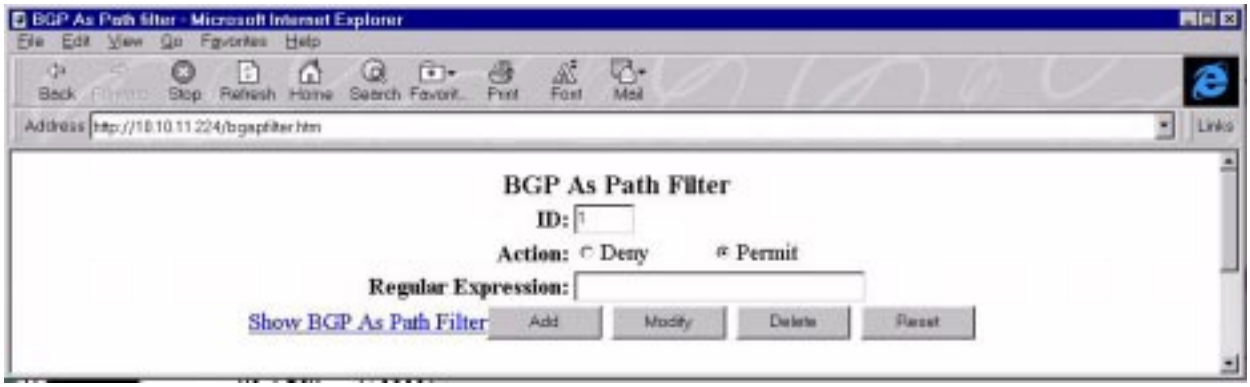

Figure 12.11 BGP AS Path filter panel

Using Regular Expressions

You use a regular expression for the <AS-path > parameter to specify a single character or multiple characters as a filter pattern. If the AS-path matches the pattern specified in the regular expression, the filter evaluation is true; otherwise, the evaluation is false.

In addition, you can include special characters that influence the way the software matches the AS-path against the filter value. For example, you can enter a list of characters followed by "**.**" (a period) to cause the filter comparison to be true for any AS-path that contains at least one of the characters in the regular expression.

To filter on a specific single-character value, enter the character for the <AS-path> parameter. For example, to filter on AS-paths that contain the letter "z", enter the following command:

```
HP9300(config-bgp-router)# as-path-filter 1 permit z
```

To filter on a string of multiple characters, enter the characters in brackets. For example, to filter on AS-paths that contain "x", "y", or "z", enter the following command:

```
HP9300(config-bgp-router)# as-path-filter 1 permit [xyz]
```

Special Characters

When you enter as single-character expression or a list of characters, you also can use the following special characters. Table 12.1 lists the special characters. The description for each special character includes an example. Notice that you place some special characters in front of the characters they control but you place other special characters after the characters they control. In each case, the examples show where to place the special character.

Table 12.1: BGP4 Special Characters for Regular Expressions

Character	Operation
.	The period matches on any single character within a list of characters, including a blank space. For example, the following regular expression matches for "d", "e", or "g": deg.
*****	The asterisk matches on zero or more sequences of a pattern. For example, the following regular expression matches on an AS-path that contains the string "1111" followed by any value: 1111*

Table 12.1: BGP4 Special Characters for Regular Expressions (Continued)

Character	Operation
+	The plus sign matches on one or more sequences of a pattern. For example, the following regular expression matches on an AS-path that contains the sequence "deg": deg+
?	The question mark matches on zero occurrences or one occurrence of a pattern. For example, the following regular expression matches on an AS-path that contains "dg" or "deg": de?g
^	The caret matches on any characters *except* the ones in the regular expression. For example, the following regular expression matches on an AS-path that does *not* contain "deg": ^deg
$	A dollar sign matches on the end of an input string. For example, the following regular expression matches on an AS-path that ends with "deg": $deg
_	An underscore matches on any of the following: • **,** (comma) • **{** (left curly brace) • **}** (right curly brace) • **(** (left parenthesis) • **)** (right parenthesis) • The beginning of the input string • The end of the input string • A blank space For example, the following regular expression matches on an AS-path that ends with a period: $_
-	A hyphen separates the beginning and ending of a range of characters. For example, the following regular expression matches on an AS-path that contains "abcdefghijklm": [a-m]
[ ]	Square brackets enclose a range of single-character patterns. For example, the following regular expression matches on an AS-path that contains "12345": [1-5]
|	A vertical bar (sometimes called a pipe or a logical "or") separates two alternative values or sets of values. The AS-path can match one or the other value. For example, the following regular expression matches on an AS-path that contains either "12345" or "6789": [1-5]|[6-9]

If you want to filter for a special character instead of using the special character as described in Table 12.1, enter "****" (backslash) in front of the character. For example, to filter on AS-path strings containing an asterisk, enter the asterisk portion of the regular expression as "*****".

```
HP9300(config-bgp-router)# as-path-filter 2 deny \*
```

If you want to filter on multiple instances of the same character or pattern or characters within an AS-path, you can use parentheses followed by "\<num>", where <num> causes the pattern to be reused later in the regular expression. For example, to filter on multiple instances of the pattern "zyx" in an AS-path, use the following command:

```
HP9300(config-bgp-router)# as-path-filter 2 deny zyx\1
```

This command creates AS-path filter 2 to filter AS-paths for multiple instances of the pattern "zyx".

Filtering Communities

A community is an optional attribute that identifies the route as a member of a user-defined class of routes. Community names are arbitrary numerical strings up to eight characters long. You determine what the name means when you create the community name as one of a route's attributes. You can define up to 100 community filters for BGP4.

NOTE: Once you define a filter, the default action for communities that do not match a filter is "deny". To change the default action to "permit", configure the last filter as "permit any any".

Community filters can be referred to by match statements in a route map.

You can filter routes received from BGP4 neighbors based on community names. Use either of the following methods to do so.

USING THE CLI

To define filter 3 to permit routes that have the NO_ADVERTISE community, enter the following command:

```
HP9300(config-bgp-router)# community-filter 3 permit no-advertise
```

syntax: community-filter <num> permit|deny <community-number>|internet|no-advertise|no-export

The <num> parameter identifies the filter's position in the community filter list and can be from 1 – 100. Thus, the community filter list can contain up to 100 filters. The routing switch applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the routing switch stops and does not continue applying filters from the list.

NOTE: If the filter is referred to by a route map's match statement, the filter is applied in the order in which the filter is listed in the match statement.

The **permit|deny** parameter indicates the action the routing switch takes if the filter match is true.

- If you specify **permit**, the routing switch permits the route into the BGP4 table if the filter match is true.
- If you specify **deny**, the routing switch denies the route from entering the BGP4 table if the filter match is true.

The <community-number> parameter indicates a specific community number to filter. Use this parameter to filter for a private (administrator-defined) community. If you want to filter for the well-known communities "NO_EXPORT" or "NO_ADVERTISE", use the corresponding keyword (described below).

The **internet** keyword checks for routes that do not have the community attribute. Routes without a specific community are considered by default to be members of the largest community, the Internet.

The **no-advertise** keyword filters for routes with the well-known community "NO_ADVERTISE". A route in this community should not be advertised to any BGP4 neighbors.

The **no-export** keyword filters for routes with the well-known community "NO_EXPORT". A route in this community should not be advertised to any BGP4 neighbors outside the local AS.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Community Filter link from the bottom of the BGP configuration sheet to display the BGP Community Filter panel, shown in Figure 12.12.
2. Enter the filter ID in the ID field. You can specify a number from 1 – 100.
3. Select the action you want the routing switch to perform if the filter is true.
   - If you select Deny, the routing switch denies the route from entering the BGP4 table if the filter match is true.
   - If you select Permit, the routing switch permits the route into the BGP4 table if the filter match is true.
4. If you are filtering on a well-known community type, select the type:
   - Internet – Filters for routes that do not have the community attribute. Routes without a specific community are considered by default to be members of the largest community, the Internet.
   - No Advertise – Filters filters for routes with the well-known community "NO_ADVERTISE". A route in this community should not be advertised to any BGP4 neighbors.
   - No Export – Filters for routes with the well-known community "NO_EXPORT". A route in this community should not be advertised to any BGP4 neighbors outside the local AS.

   NOTE: If you want to filter on a private (administrator-defined) community, do not select on of these. Instead, enter the community number in the Community List field.
5. If you are filtering on a private community, enter the community number in the Community List field.
6. Select Add to add the filter.

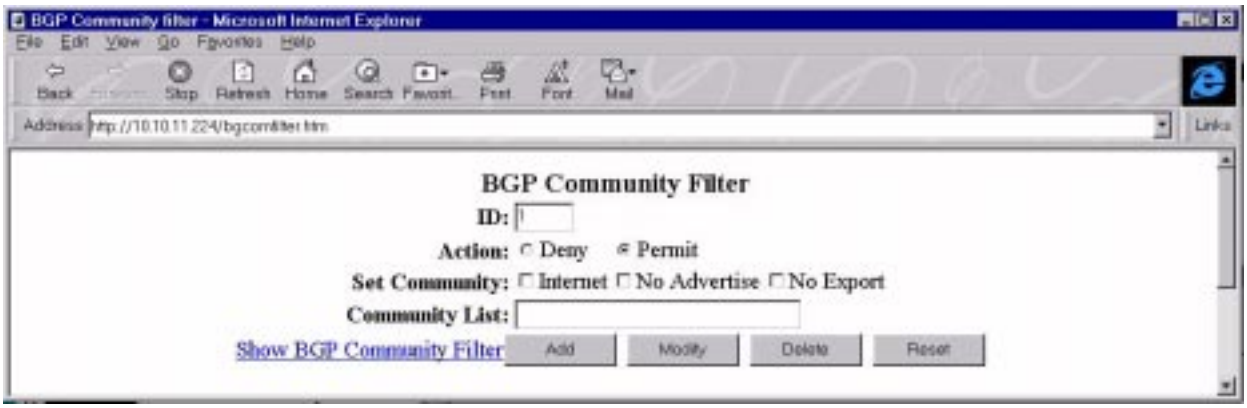

Figure 12.12 BGP Community Filter link

Defining Route Maps

A ***route map*** is a named set of match conditions and parameter settings that the routing switch can use to modify route attributes and to control redistribution of the routes into other protocols. A route map consists of a sequence of up to 50 ***instances***. If you think of a route map as a table, an instance is a row in that table. The routing switch evaluates a route according to a route map's instances in ascending numerical order. The route is first compared against instance 1, then against instance 2, and so on. As soon as a match is found, the routing switch stops evaluating the route against the route map instances.

Route maps can contain ***match*** statements and ***set*** statements. Each route map contains a “permit” or “deny” action for routes that match the match statements.

- If the route map contains a permit action, a route that matches a match statement is permitted; otherwise, the route is denied.
- If the route map contains a deny action, a route that matches a match statement is denied.
- If a route does not match any match statements in the route map, the route is denied. This is the default action. To change the default action, configure the last match statement in the last instance of the route map to "permit any any".
- For route maps that contain address filters, AS-path filters, or community filters, if the action specified by a filter conflicts with the action specified by the route map, the route map’s action takes precedence over the individual filter’s action.

NOTE: You can configure up to 50 route maps. Regardless of the number of route maps you configure, you can have a combined total of 300 route map instances (sequences), match statements, and set statements.

If the route map contains set statements, routes that are permitted by the route map’s match statements are modified according to the set statements.

Match statements compare the route against one or more of the following:

- The route’s BGP4 MED (metric)
- A sequence of AS-path filters
- A sequence of community filters
- A sequence of address filters
- The IP address of the next hop router
- The route’s tag
- For OSPF routes only, the route’s type (internal, external type-1, or external type-2)

For routes that match at least one of the match statements, the route map’s set statements can perform one or more of the following modifications to the route’s attributes:

- Prepend AS numbers to the front of the route’s AS-path. By adding AS numbers to the AS-path, you can cause the route to be less preferred when compared to other routes on the basis of the length of the AS-path.
- Add a user-defined tag to the route or add an automatically calculated tag to the route.
- Set the community value to the well-known value “NO_EXPORT” or “NO_ADVERTISE”, or set it to “none”.

  NOTE: The community "none" is equivalent to the community "internet", which can be checked for using a community filter. See “Filtering Communities” on page 12-35.

- Set the local preference.
- Set the MED (metric).
- Set the IP address of the next hop router.
- Set the origin to IGP or INCOMPLETE.
- Set the weight.

For example, when you configure parameters for redistributing routes into RIP, one of the optional parameters is a route map. If you specify a route map as one of the redistribution parameters, the routing switch will match the route against the match statements in the route map. If a match is found and if the route map contains set statements, the routing switch will set attributes in the route according to the set statements.

To create a route map, you define instances of the map. Each instance is identified by a sequence number. A route map can contain up to 50 instances.

To define a route map, use the procedures in the following sections.

Entering the Route Map Into the Software

USING THE CLI

To add instance 1 of a route map named "GET_ONE" with a permit action, enter the following command.

```
HP9300(config)# route-map GET_ONE permit 1
HP9300(config-bgp-routemap GET_ONE)#
```

syntax: route-map <map-name> permit|deny <num>

As shown in this example, the command prompt changes to the Route Map level. You can enter the match and set statements at this level. See “Specifying the Match Conditions” on page 12-39 and “Setting Parameters in the Routes” on page 12-40.

The <map-name> is a string of characters that names the map. Map names can be up to 32 characters in length. You can define up 50 route maps on the routing switch.

The **permit|deny** parameter specifies the action the routing switch will take if a route matches a match statement.

- If you specify **deny**, the routing switch does not advertise or learn the route.
- If you specify **permit**, the routing switch applies the match and set statements associated with this route map instance.

The <num> parameter specifies the instance of the route map you are defining. Each route map can have up to 50 instances.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Route Map Filter link at the bottom of the BGP configuration sheet to display the BGP Route Map Filter panel, shown in Figure 12.13.
2. Enter the name of the route map in the Route Map Name field.
3. Enter the sequence (instance) number in the Sequence field. The routing switch applies the instances in ascending numerical order. Once an instance comparison results in a "true" evaluation, the routing switch stops applying instances and applies the match and set statements you configure for the instance. See “Specifying the Match Conditions” on page 12-39 and “Setting Parameters in the Routes” on page 12-40.
4. Select the action you want the routing switch to perform if the comparison results in a "true" value.
   - If you select Deny, the routing switch does not advertise or learn the route.
   - If you select Permit, the routing switch applies the match and set statements associated with this route map instance.
5. Select Add to add the filter.

.

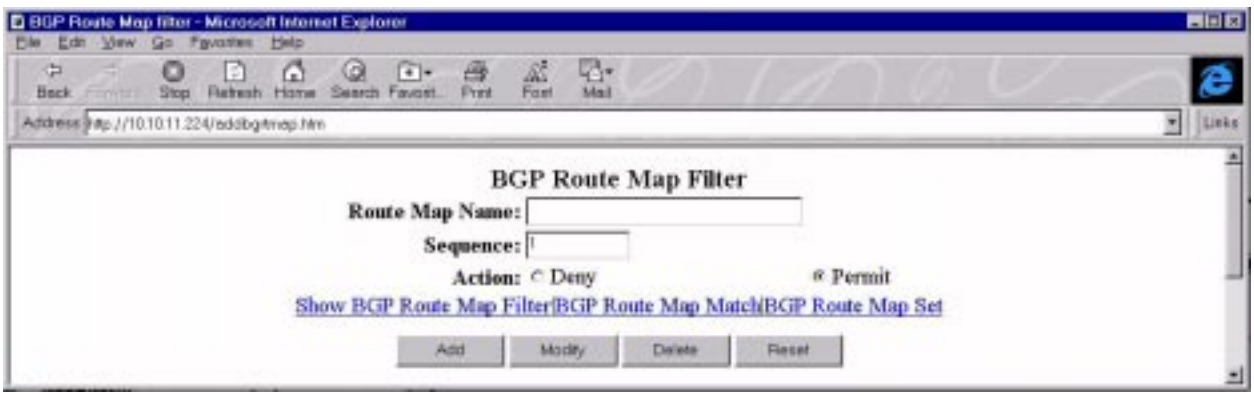

Figure 12.13 BGP Route Map Filter panel

Specifying the Match Conditions

Use the following command to define the match conditions for instance 100 of the route map GET_ONE. This instance compares the route updates against BGP4 address filter 11.

```
HP9300(config-bgp-routemap GET_ONE)# match address-filters 11
```

syntax: match [as-path-filters|community-filters|address-filters <num,num,...>] | [metric <num>] | [next-hop <address-filter-list>] | [route-type internal|external-type1|external-type2] | [tag <tag-value>]

The **as-path-filters|community-filters|address-filters** <num,num,...> specifies a filter or list of filters to be matched for each route. The routing switch treats the first match as the best match. If a route does not match any filter in the list, then the routing switch considers the match condition to have failed.

NOTE: The filters must already be configured.

The **metric** <num> parameter compares the route's MED (metric) to the specified value.

The **next-hop** <address-filter-list> parameter compares the IP address of the route's next hop to the specified IP address filters. The filters must already be configured.

The **route-type internal|external-type1|external-type2** parameter applies only to OSPF routes. This parameter compares the route's type to the specified value.

The **tag** <tag-value> parameter compares the route's tag to the specified value.

USING THE WEB MANAGEMENT INTERFACE

NOTE: To simplify testing and configuration, you can specify an option and then choose whether to activate it. To activate an option, select the checkbox in front of the option's field. Leave the checkbox unselected to leave the option inactive.

1. Select the BGP Route Map Match link from the bottom of the BGP Route Map Filter panel. A panel such as the one shown in Figure 12.14 is displayed.

   NOTE: If you have already configured some match statements for the route map, a table listing the match statements is displayed. In this case, select BGP Route Map Match at the bottom of the panel.

2. Select the sequence (instance) from the Route Map Name Sequence field's pulldown list. The routing switch applies the instances in ascending numerical order and stops after the first match.
3. For OSPF routes, select the one of the following route types—Internal, External1, or External2.
4. Optionally select the AS Path Filter checkbox and list a sequence of AS-path filters in the following field. This instance of the route map will compare the AS-path in BGP4 route updates against the AS-path filters in the order you list them in this field. The same is true for community filters and address filters.

   NOTE: The AS-path, community, and address filters must already be configured.

5. Optionally select the Community Filter checkbox and list a sequence of community filters in the following field.
6. Optionally select the Address Filter checkbox and list a sequence of community filters in the following field.
7. Optionally enter an IP address against which you want to compare the route updates' next-hop attribute. Enter the address in the Next Hop List field. Also select the checkbox in front of the field.
8. Optionally enter a tag value against which you want to compare the updates in the Tag List field. Also select the checkbox in front of the field.
9. Optionally enter a MED (metric) value against which you want to compare the route updates in the Metric field. Also select the checkbox in front of the field.
10. Select Apply to implement the changes.

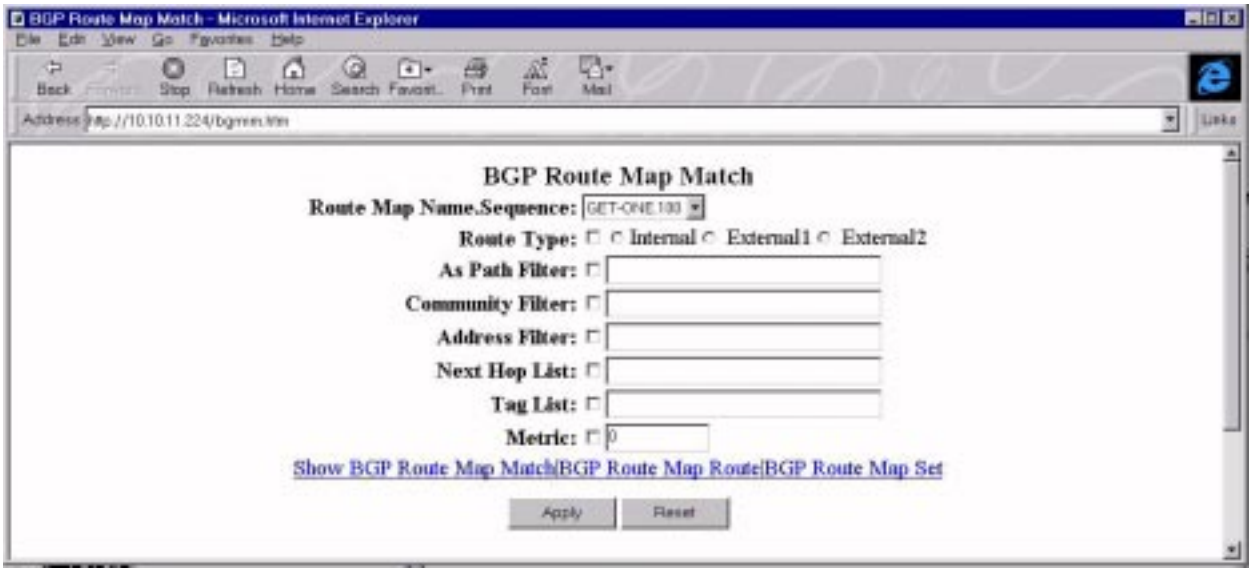

Figure 12.14 BGP Route Map Match panel

Setting Parameters in the Routes

Use the following command to define a set statement that prepends an AS number to the AS path on each route that matches the corresponding match statement.

```
HP9300(config-bgp-routemap GET_ONE)# set as-path prepend 65535
```

syntax: set as-path [prepend <AS-num,AS-num,...>] | [automatic-tag] | [community no-export|no-advertise|none] | [local-preference <num>] | [metric <num>] | [next-hop <IP-addr>] | [origin igp|incomplete] | [tag <tag-value>] | [weight <num>]

The **as-path prepend** <num,num,...> parameter adds the specified AS numbers to the front of the AS-path list for the route.

The **automatic-tag** parameter calculates and sets an automatic tag value for the route.

NOTE: This parameter applies only to routes redistributed into OSPF.

The **community no-export|no-advertise|none** parameter sets the community attribute for the route to “NO_EXPORT”, “NO_ADVERTISE”, or “none”.

The **local-preference** <num> parameter sets the community attribute for the route. The default local preference is 100. You can set the preference to a value from 0 – 4294967295.

The **metric** <num> parameter sets the MED (metric) value for the route. The default MED value is 0. You can set the preference to a value from 0 – 4294967295.

The **next-hop** <IP-addr> parameter sets the IP address of the route’s next hop router.

The **origin igp|incomplete** parameter sets the route’s origin to IGP or INCOMPLETE.

The **tag** <tag-value> parameter sets the route’s tag. You can specify a tag value from 0 – 4294967295.

NOTE: This parameter applies only to routes redistributed into OSPF.

NOTE: You also can set the tag value using a table map. The table map changes the value only when the routing switch places the route in the IP route table instead of changing the value in the BGP route table. See "Using a Table Map To Set the Tag Value" on page 12-42.

The **weight** <num> parameter sets the weight for the route. You can specify a weight value from 0 – 4294967295.

USING THE WEB MANAGEMENT INTERFACE

NOTE: To simplify testing and configuration, you can specify an option and then choose whether to activate it. To activate an option, select the checkbox in front of the option's field. Leave the checkbox unselected to leave the option inactive.

1. Select the BGP Route Map Set link from the bottom of the BGP Route Map Filter panel. A panel such as the one shown in Figure 12.15 is displayed.

   NOTE: If you have already configured some set statements for the route map, a table listing the set statements is displayed. In this case, select BGP Route Map Set at the bottom of the panel.

2. Select the sequence (instance) from the Route Map Name Sequence field's pulldown list.
3. Optionally select the origin. You can select IGP or Incomplete. Also select the checkbox in front of the field.
4. Optionally enter AS numbers to append to the AS path. Also select the checkbox in front of the field.
5. Optionally select Auto Tag. The routing switch calculates and sets an automatic tag value for the route.
6. If you did not select Auto Tag and you instead want to set the tag value manually, enter a tag value from 0 – 4294967295 in the Tag field. Also select the checkbox in front of the field.
7. Optionally select the community type and also select the checkbox.
8. For a private community, enter the community number in the Number field.
9. Select Additive Community of you want the Set statement to add the specified community.
10. Optionally enter a local preference in the Local Preference and also select the checkbox in front of the field. The default local preference is 100. You can set the preference to a value from 0 – 4294967295.
11. Optionally enter a metric (MED) in the Metric field and also select the checkbox in front of the field. The default MED value is 0. You can set the preference to a value from 0 – 4294967295.
12. Optionally enter the Next Hop IP address in the NextHop field and also select the checkbox in front of the field.
13. Optionally enter a weight in the Weight field and also select the checkbox in front of the field. You can specify a weight value from 0 – 4294967295.
14. Select Apply to implement the changes.

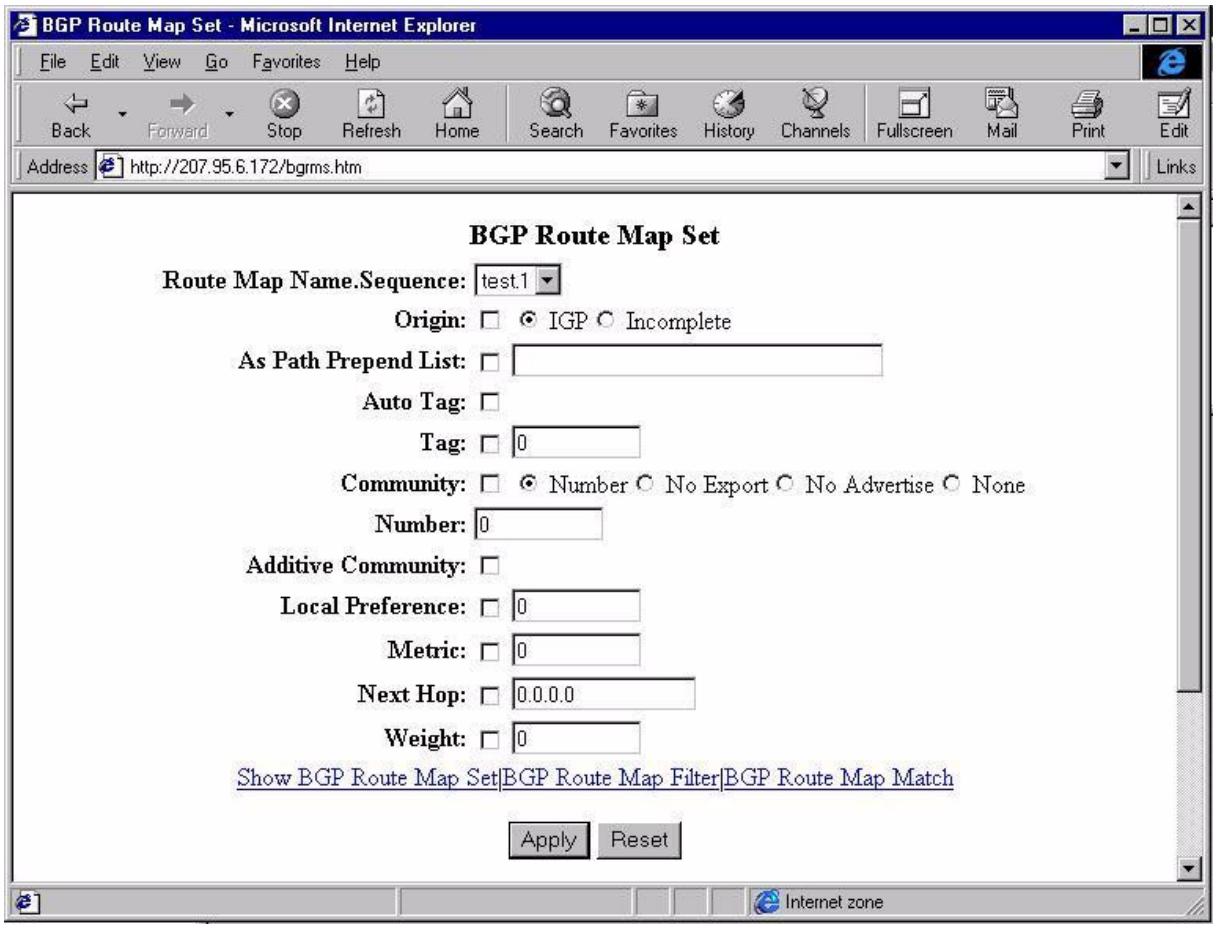

Figure 12.15 BGP Route Map Set panel

Using a Table Map To Set the Tag Value

Route maps that contain set statements change values in routes when the routes are filtered by the route map. For inbound route maps (route maps that filter routes received from neighbors), this means that the routes are changed before they enter the BGP4 route table.

For tag values, if you do not want the value to change until a route enters the IP route table, you can use a table map to change the value. A table map is a route map that you have associated with the IP routing table. The routing switch applies the set statements for tag values in the table map to routes before adding them to the route table.

To configure a table map, you configure the route map, then identify it as a table map. The table map does not require separate configuration. You create it simply by calling an existing route map a table map. You can have one table map.

NOTE: Use table maps only for setting the tag value. Do not use table maps to set other attributes. To set other route attributes, use route maps or filters.

USING THE CLI

To create a route map and identify it as a table map, enter commands such as following. These commands create a route map that uses an address filter. For routes that match the address filter, the route map changes the tag value to 100. This route map is then identified as a table map. As a result, the route map is applied only to routes that the routing switch places in the IP route table. The route map is not applied to all routes. This example assumes that address filter 11 has already been configured.

```
HP9300(config)# route-map TAG_IP permit 1
HP9300(config-bgp-routemap TAG_IP)# match address-filters 11
HP9300(config-bgp-routemap TAG_IP)# set tag 100
HP9300(config-bgp-routemap TAG_IP)# router bgp
HP9300(config-bgp-router)# table-map TAG_IP
```

USING THE WEB MANAGEMENT INTERFACE

1. Use the Web management procedures in “Defining Route Maps” on page 12-36 to create the route map.
2. Select the BGP link to display the panel shown in Figure 12.3.
3. Select the route map name from the Map field’s pulldown menu.
4. Select Apply to implement the change.

Displaying BGP4 Information

You can display the following configuration information and statistics for the BGP4 protocol on the routing switch:

- Summary BGP4 configuration information for the routing switch
- Information about the routing switch’s BGP4 neighbors
- Information about the paths from which BGP4 selects routes
- The routing switch’s BGP4 route table

Displaying Summary BGP4 Information

You can display the local AS number, the maximum number of routes and neighbors supported, and some BGP4 statistics using either of the following methods.

USING THE CLI

To view summary BGP4 information for the routing switch, enter the following command at any CLI prompt:

```
HP9300# show ip bgp summary
```

Here is an example of the information displayed by this command:

```
HP9300# show ip bgp summary
  BGP4 Summary
  Local AS Number : 65002
  Maximum Number of Attribute Entries Supported :10000
  Maximum Number of Routes Supported : 60000
  Maximum Number of Neighbors Supported : 3
  Number of Routes Installed : 58756
  Number of Attribute Entries Installed : 7750
  Neighbor Address  State         StateChangeTime  RoutesRecvd  RoutesInstalled
  192.168.11.1      ESTABLISHED   0  :0  :43 :54   58737        58737
  192.168.88.28     ESTABLISHED   0  :2  :26 :43   0            0
  192.168.199.1     ESTABLISHED   0  :0  :48 :5    13           13
```

This display shows the following information.

Table 12.2: BGP4 Summary Information

This Field...	Displays...
Local AS Number	The BGP4 AS number the routing switch is in.
Maximum Number of Attribute Entries Supported	The number of attribute entries the routing switch's memory can hold. An attribute entry is a set of route attributes that are associated with one or more routes. To reconfigure the memory size for entries, see "Changing the Maximum Number of Route-Attribute Entries" on page 12-18.
Maximum Number of Routes Supported	The number of BGP4 routes the routing switch's memory can hold. To reconfigure the memory size for routes, see "Changing the Maximum Number of Routes" on page 12-18.
Maximum Number of Neighbors Supported	The number of BGP4 neighbors the routing switch can have. To reconfigure the memory size for neighbors, see "Changing the Maximum Number of Neighbors" on page 12-17.
Number of Routes Installed	The number of BGP4 routes in the routing switch's BGP4 route table. To display the BGP4 route table, see "Displaying the BGP4 Route Table" on page 12-52.
Number of Attribute Entries Installed	The number of BGP4 route-attribute entries in the routing switch's route-attributes table. To display the route-attribute table, see "Displaying BGP4 Route-Attribute Entries" on page 12-55.
Neighbor Address	The IP addresses of this routing switch's BGP4 neighbors.

Table 12.2: BGP4 Summary Information (Continued)

This Field...	Displays...
State	The state of this routing switch's neighbor session with each neighbor. The states are from this routing switch's perspective of the session, not the neighbor's perspective. The state values are based on the BGP4 state machine values described in RFC 1771 and can be one of the following for each routing switch: • IDLE – The BGP4 process is waiting to be started. Usually, enabling BGP4 or establishing a neighbor session starts the BGP4 process. • CONNECT – BGP4 is waiting for the connection process for the TCP neighbor session to be completed. • ACTIVE – BGP4 is trying to open a TCP connection to the neighbor. **Note**: If the state frequently changes between CONNECT and ACTIVE, there may be a problem with the TCP connection. • OPEN SENT – BGP4 is waiting for an Open message from the neighbor. • OPEN CONFIRM – BGP4 has received an OPEN message from the neighbor and is now waiting for either a KEEPALIVE or NOTIFICATION message. If the routing switch receives a KEEPALIVE message from the neighbor, the state changes to Established. If the message is a NOTIFICATION, the state changes to Idle. • ESTABLISHED – BGP4 is ready to exchange UPDATE packets with the neighbor.
StateChangeTime	The time that has passed since the state last changed. The time is shown in days:hours:minutes:seconds.
RoutesRcvd	The total number or routes received in UPDATE messages from the neighbor since the session was first established.
RoutesInstalled	The number of routes received from the neighbor that this routing switch actually installed in the BGP4 route table. Usually, this number is lower than the RoutesRcvd number. The difference indicates that this routing switch filtered out some of the routes received in the UPDATE messages.

USING THE WEB MANAGEMENT INTERFACE

1. Select Show.
2. Select Summary in the BGP section. A table such as the one shown in Figure 12.16 is displayed.

Figure 12.16 BGP Summary display

Displaying BGP4 Neighbor Information

You can display configuration information and statistic for the routing switch's BGP4 neighbors using either of the following methods.

USING THE CLI

To view BGP4 neighbor information for the routing switch, enter the following command:

```
HP9300# show ip bgp neighbors
```

syntax: show ip bgp neighbor [<IP-addr> [advertised-routes] [last-packet-with-error] [attribute-entries] [received-routes] [routes-summary]]

The <IP-addr> option lets you narrow the scope of the command to a specific neighbor.

The **advertised-routes** option displays only the routes that the routing switch has advertised to the neighbor during the current BGP4 neighbor session.

The **last-packet-with-error** displays a hexadecimal dump of the first 400 bytes of the last packet received from the neighbor that contained an error.

The **attribute-entries** option shows the attribute-entries associated with routes received from the neighbor.

The **received-routes** option lists the routes received in UPDATE messages from the neighbor.

The **routes-summary** option displays a summary of the following information:

- Number of routes received from the neighbor
- Number of routes accepted by this routing switch from the neighbor
- Number of routes this routing switch filtered out of the UPDATES received from the neighbor and did not accept
- Number of routes advertised to the neighbor
- Number of attribute entries associated with routes received from or advertised to the neighbor.

Here is an example of the information displayed by this command. In this example, a specific neighbor's IP address is entered. The command therefore shows information only for that neighbor. None of the other options are used; thus, all the information about the neighbor is displayed. The numbers in the leftmost column separate the entries for each neighbor.

```
HP9300# show ip bgp neighbors 192.168.11.1
        Total number of BGP Neighbors: 3
        IP Address           Remote AS            EBGP/IBGP        State
1       192.168.11.1         65001                EBGP             ESTABLISHED
        Keep Alive Time      Hold Time            Advertisement Interval
        0                    0                    5
                             Message Sent         Message Received
        Keep Alive           3                    3
        Update               19                   28270
        Notifications        0                    0
        Open                 3                    3
        Last Connection Reset Reason:Port State Down
        Notification Message Error Code Received:Unspecified
        Notification Message Error SubCode Received:Not Applicable
        Notification Message Error Code Transmitted:Unspecified
        Notification Message Error SubCode Transmitted:Not Applicable
```

This example shows how to display information for a specific neighbor, by specifying the neighbor's IP address with the command. None of the other display options are used; thus, all of the information is displayed for the neighbor. The number in the far left column indicates the neighbor for which information is displayed. When you list information for multiple neighbors, this number makes the display easier to read.

This display shows the following information.

Table 12.3: BGP4 Neighbor Information

This Field...	Displays...
Total number of BGP Neighbors	The total number of neighbors configured on this routing switch.
IP Address	The IP address of the neighbor.
Remote AS	The AS the neighbor is in.
EBGP/IBGP	Whether the neighbor session is an IBGP session or an EBGP session. If the neighbor is in the same AS, the session is IBGP. If the neighbor is in another AS, the session is EBGP.

Table 12.3: BGP4 Neighbor Information (Continued)

This Field...	Displays...
State	The state of the routing switch's session with the neighbor. The states are from this routing switch's perspective of the session, not the neighbor's perspective. The state values are based on the BGP4 state machine values described in RFC 1771 and can be one of the following for each routing switch: • IDLE – The BGP4 process is waiting to be started. Usually, enabling BGP4 or establishing a neighbor session starts the BGP4 process. • CONNECT – BGP4 is waiting for the connection process for the TCP neighbor session to be completed. • ACTIVE – BGP4 is trying to open a TCP connection to the neighbor. **Note**: If the state frequently changes between CONNECT and ACTIVE, there may be a problem with the TCP connection. • OPEN SENT – BGP4 is waiting for an Open message from the neighbor. • OPEN CONFIRM – BGP4 has received an OPEN message from the neighbor and is now waiting for either a KEEPALIVE or NOTIFICATION message. If the routing switch receives a KEEPALIVE message from the neighbor, the state changes to Established. If the message is a NOTIFICATION, the state changes to Idle. • ESTABLISHED – BGP4 is ready to exchange UPDATE messages with the neighbor.
Keep Alive Time	The keep alive time, which specifies how often this routing switch sends keep alive messages to the neighbor. See "Changing the Keep Alive Time and Hold Time" on page 12-16.
Hold Time	The hold time, which specifies how many seconds the routing switch will wait for a KEEPALIVE or UPDATE message from a BGP4 neighbor before deciding that the neighbor is dead. See "Changing the Keep Alive Time and Hold Time" on page 12-16.

Table 12.3: BGP4 Neighbor Information (Continued)

This Field...	Displays...
Advertisement Interval	The advertisement interval, which specifies the minimum delay (in seconds) between messages to the neighbor. The default is 30 for EBGP neighbors (neighbors in other ASs). The default is 5 for IBGP neighbors (neighbors in the same AS). The range is 0 – 600. To change the advertisement interval, see “Adding BGP4 Neighbors” on page 12-13. **Note**: The routing switch applies the advertisement interval only under certain conditions. The routing switch does not apply the advertisement interval when sending initial updates to a BGP4 neighbor. As a result, when a routing switch needs to send its entire routing table to a BGP4 neighbor, the routing switch sends the updates one immediately after another at rate of one TCP window per second, without waiting for the advertisement interval. The routing switch still applies the advertisement interval to an update if the update contains a route for which the routing switch has just sent an update. For example, if the routing switch sends an update for routes 1,2, and 3, then receives a change to an attribute of one of the routes before the advertisement interval has expired, the routing switch waits to send an update for the change until the advertisement interval has expired.
Message Sent	The number of messages this routing switch has sent to the neighbor during the current BGP4 session. This counter is reset to zero if the session ends.
Message Received	The number of messages this routing switch has received from the neighbor during the current BGP4 session. This counter is reset to zero if the session ends.
Keep Alive	For the Message Sent and Message Received columns, indicates how many KEEPALIVE messages have been sent and received by this routing switch.
Update	For the Message Sent and Message Received columns, indicates how many UPDATE messages have been sent and received by this routing switch.
Notifications	For the Message Sent and Message Received columns, indicates how many NOTIFICATION messages have been sent and received by this routing switch.
Open	For the Message Sent and Message Received columns, indicates how many OPEN messages have been sent and received by this routing switch.

Table 12.3: BGP4 Neighbor Information (Continued)

This Field...	Displays...
Last Connection Reset Reason	The reason the previous session with this neighbor ended. The reason can be one of the following: • Update Msg with Header Error • Open Msg with Error • Hold Timer Expired • User Closed Peer Session • TCP Connection Closed • Update Msg with Bad Attribute List • Update Msg with Unrecognized Attribute • Update Msg with Missing Attribute • Update Msg with Attribute Error • Update Msg with Attribute Length Error • Update Msg with Bad Origin Attribute • As Path Loop • Invalid Next Hop Attribute • Update Msg with Optional Attribute Error • Invalid Network Field • Update Msg with Bad AS Path Attribute • Notification Error Message Received • Port State Down • Finite State Machine Error • TCP Connection Closed by Remote • Reason Unknown

Table 12.3: BGP4 Neighbor Information (Continued)

This Field...	Displays...

Notification Message Error Code Received	If the routing switch receives a NOTIFICATION messages from the neighbor, the message contains an error code corresponding to one of the following errors. Some errors have subcodes that clarify the reason for the error. Where applicable, the subcode messages are listed underneath the error code messages.
	• Message Header Error
	  • Connection Not Synchronized
	  • Bad Message Length
	  • Bad Message Type
	  • Unspecified
	• Open Message Error
	  • Unsupported Version
	  • Bad Peer As
	  • Bad BGP Identifier
	  • Unsupported Optional Parameter
	  • Authentication Failure
	  • Unacceptable Hold Time
	  • Unspecified
	• Update Message Error
	  • Malformed Attribute List
	  • Unrecognized Attribute
	  • Missing Attribute
	  • Attribute Flag Error
	  • Attribute Length Error
	  • Invalid Origin Attribute
	  • AS Routing Loop
	  • Invalid NextHop Attribute
	  • Optional Attribute Error
	  • Invalid Network Field
	  • Malformed AS Path
	  • Unspecified
	• Hold Timer Expired
	• Finite State Machine Error
	• Cease
	• Unspecified

Notification Message Error SubCode Received	See above.

Table 12.3: BGP4 Neighbor Information (Continued)

This Field...	Displays...
Notification Message Error Code Transmitted	The error message corresponding to the error code in the NOTIFICATION message this routing switch sent to the neighbor. See the description for the Notification Message Error Code Received field for a list of possible codes.
Notification Message Error SubCode Transmitted	See above.

USING THE WEB MANAGEMENT INTERFACE

Select the Neighbor link from the BGP configuration sheet to display a panel such as the one shown in Figure 12.17.

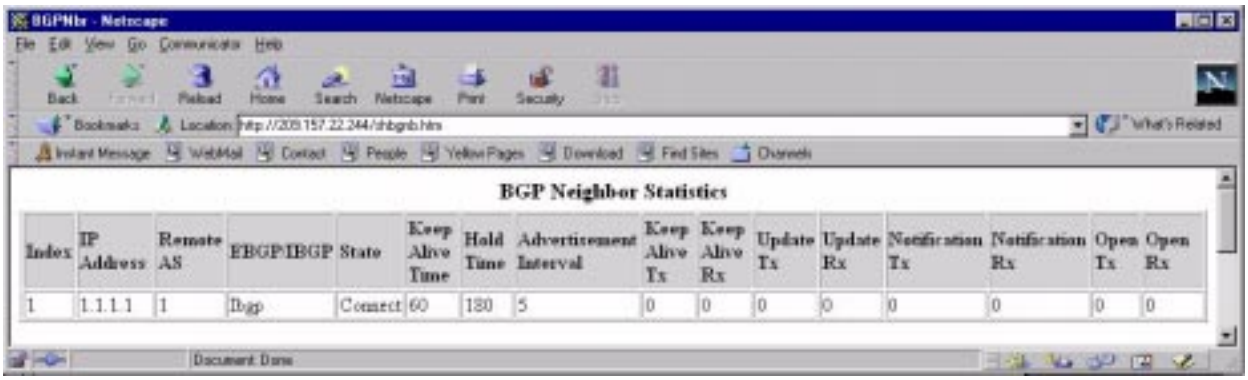

Figure 12.17 BGP Neighbors information

Displaying the BGP4 Route Table

BGP4 uses filters you define as well as the algorithm described in “How BGP4 Selects a Path for a Route” on page 12-3 to determine the preferred route to a destination. BGP4 sends only the preferred route to the routing switch’s IP table. However, if you want to view all the routes BGP4 knows about, you can display the BGP4 table using either of the following methods.

USING THE CLI

To view the BGP4 route table, enter the following command:

```
HP9300# show ip bgp routes
```

Syntax: show ip bgp routes <num> [cidr-only] [community <num>|no-export|no-advertise|internet] [community-list <num>] [detail <option>] [filter-list <num, num,...>] [network <IP-addr>] [regular-expression <value>]

The <num> option specifies the table entry with which you want the display to start. For example, if you want to list entries beginning with table entry 100, specify 100.

The **cidr-only** option lists only the routes that do not have a mask length of 8, 16, or 24bits (the standard Class A, B, and C sub-net mask lengths).

The **community** option lets you displays routes for a specific community. You can specify **no-export**, **no-advertise**, **internet**, or a private community number.

The **community-list** option lets you display routes that match a specific community filter.

The **detail** option lets you display greater detail for one of the other options.

The **filter-list** option displays routes that match a specific address filter list.

The **network** option displays routes for a specific network.

The **regular-expression** option filters the display based on a regular expression. See “Using Regular Expressions” on page 12-33.

Here is an example of the information displayed by this command:

```
HP9300#show ip bgp routes
       Total number of BGP Routes: 58788
       Status A: AGGREGATE B:BEST  I:INTERNAL L: LOCAL  S: SUPPRESSED
       Network         ML Next Hop        Metric     LocPrf     Weight Status
1      8.9.253.160     27 192.168.11.1    0          100        0      B
2      12.0.0.0        8  192.168.11.1    0          100        0      B
3      12.2.97.0       24 192.168.11.1    0          100        0      B
4      12.2.169.0      24 192.168.11.1    0          100        0      B
5      12.3.123.0      24 192.168.11.1    0          100        0      B
6      12.3.63.0       24 192.168.11.1    0          100        0      B
7      12.2.109.0      24 192.168.11.1    0          100        0      B
8      12.4.5.0        24 192.168.11.1    0          100        0      B
remaining 58780 entries not shown...
```

Here is an example of the information displayed when you use the detail option. In this example, the information for one route is shown.

```
HP9300#show ip bgp routes detail
       Total number of BGP Routes: 388
       Status A: AGGREGATE B:BEST  I:INTERNAL L: LOCAL  S: SUPPRESSED
       Network          MaskLen Next Hop         Metric     LocPrf     Weight
1      12.2.97.0        24      192.168.11.1     0          100        0
       Originator       Atomic  AGGREGATION-ID   AS     Cluster List
       0.0.0.0          FALSE   0.0.0.0          0      None
       Origin  Status  Route Tag   Communities
       IGP     B       00000000    Internet
       AS Path : (65002) 65001 4355 2548 7018 10656
remaining 387 entries not shown...
```

This summary display (see the first example) shows the following information.

Table 12.4: BGP4 Route Information

This Field...	Displays...
Total number of BGP Routes	The number of routes contained in this routing switch's BGP4 route table.
Network	The network address of the route.
ML	The length of the CIDR network mask for the route. The number displayed in this column is the number of bits in the mask.
Next Hop	The IP address of the next hop router for this route.
Metric	The value of the route's MED attribute.
LocPrf	The degree of preference for this route relative to other routes in the local AS. When the BGP4 algorithm compares routes on the basis of local preferences, the route with the higher local preference is chosen. The preference can have a value from 0 – 4294967295. The default is 100. This parameter applies only to routes within the local AS.
Weight	The value that this routing switch associates with routes from a specific neighbor. For example, if the routing switch receives routes to the same destination from two BGP4 neighbors, the routing switch prefers the route from the neighbor with the larger weight.
Status	The route's status, which can be one of the following: • A – AGGREGATE. The route is an aggregate route for multiple networks. • B – BEST. Using the route selection algorithm, BGP4 has determined that this is the optimal route to the destination. See "How BGP4 Selects a Path for a Route" on page 12-3. • I – INTERNAL. The route was learned through IBGP and its destination is in the local AS. • L – LOCAL. The route originated on this routing switch. • S – SUPPRESSED. This route was suppressed during aggregation and thus is not advertised to neighbors.

USING THE WEB MANAGEMENT INTERFACE

1. Select Show.
2. Select Routes under BGP to display a table such as the one shown in Figure 12.18.

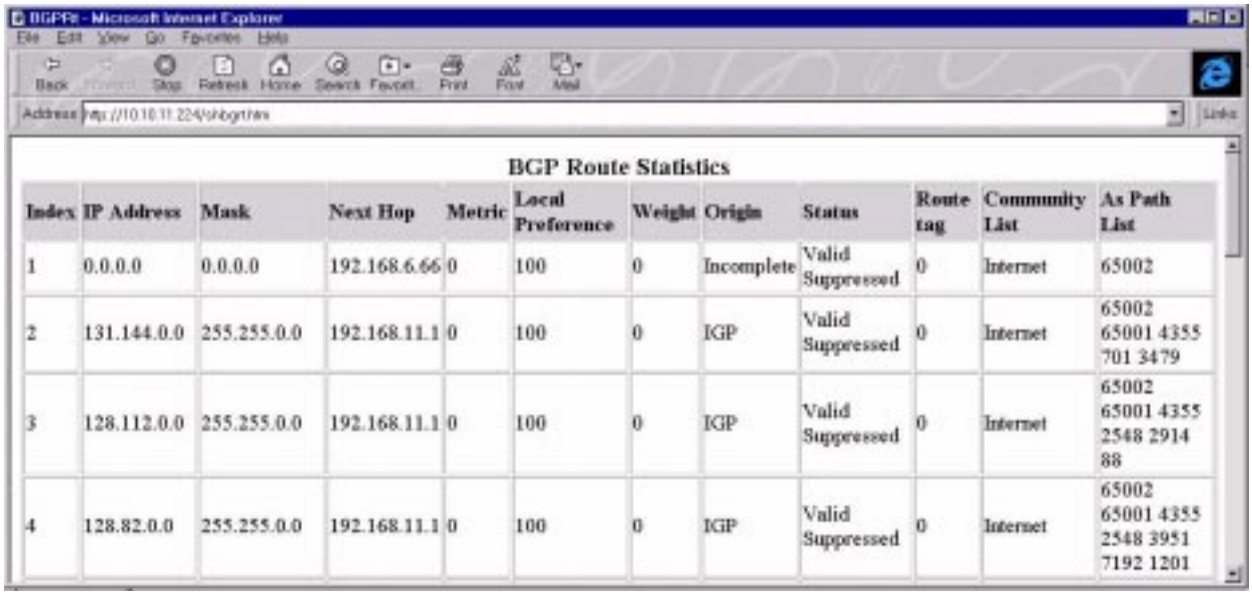

BGP Route Statistics

Index	IP Address	Mask	Next Hop	Metric	Local Preference	Weight	Origin	Status	Route tag	Community List	As Path List
1	0.0.0.0	0.0.0.0	192.168.6.66	0	100	0	Incomplete	Valid Suppressed	0	Internet	65002
2	131.144.0.0	255.255.0.0	192.168.11.1	0	100	0	IGP	Valid Suppressed	0	Internet	65002 65001 4355 701 3479
3	128.112.0.0	255.255.0.0	192.168.11.1	0	100	0	IGP	Valid Suppressed	0	Internet	65002 65001 4355 2548 2914 88
4	128.82.0.0	255.255.0.0	192.168.11.1	0	100	0	IGP	Valid Suppressed	0	Internet	65002 65001 4355 2548 3951 7192 1201

Figure 12.18 BGP Route Statistics display

Displaying BGP4 Route-Attribute Entries

The route-attribute entries table lists the sets of BGP4 attributes stored in the routing switch's memory. Each set of attributes is unique and can be associated with one or more routes. In fact, the routing switch typically has fewer route attribute entries than routes. To display the route-attribute entries table, use one of the following methods.

USING THE CLI

To display the IP route table, enter the following command:

```
HP9300# show ip attribute-entries
```

Syntax: show ip attribute-entries

Here is an example of the information displayed by this command. A zero value indicates that the attribute is not set.

```
HP9300# show ip bgp attribute-entries
        Total number of BGP Attribute Entries: 7753
1       Next Hop  :192.168.11.1       Metric   :0                Origin:IGP
        Originator:0.0.0.0            Cluster List:None
         Aggregator:AS Number :0         Router-ID:0.0.0.0          Atomic:FALSE
        Local Pref:100                Communities:Internet
        AS Path   :(65002) 65001 4355 2548 3561 5400 6669 5548
2       Next Hop  :192.168.11.1       Metric   :0                Origin:IGP
        Originator:0.0.0.0            Cluster List:None
         Aggregator:AS Number :0         Router-ID:0.0.0.0          Atomic:FALSE
        Local Pref:100                Communities:Internet
        AS Path   :(65002) 65001 4355 2548
remaining 7751 entries not shown...
```

This display shows the following information.

Table 10.5: BGP4 Route-Attribute Entries Information

This Field...	Displays...
Total number of BGP Attribute Entries	The number of routes contained in this routing switch's BGP4 route table.
Next Hop	The IP address of the next hop router for routes that have this set of attributes.
Metric	The cost of the routes that have this set of attributes.
Origin	The source of the route information. The origin can be one of the following: • EGP – The routes with this set of attributes came to this routing switch through EGP. • IGP – The routes with this set of attributes came to this routing switch through IGP. Thus, they originated in the local AS. • INCOMPLETE – The routes came from an origin other than one of the above. For example, they may have been redistributed from OSPF or RIP. When BGP4 compares multiple routes to a destination to select the best route, EGP is preferred over IGP and both are preferred over INCOMPLETE.
Originator	The route reflector that originated this set of attributes.
Cluster List	The route-reflector clusters through which this set of attributes has passed.
Aggregator	Aggregator information: • AS `Number` shows the AS in which the network information in the attribute set was aggregated. This value applies only to aggregated routes and is otherwise 0. • `Router-ID` shows the router that originated this aggregator.
Atomic	Whether the network information in this set of attributes has been aggregated *and* this aggregation has resulted in information loss. • TRUE – Indicates information loss has occurred • FALSE – Indicates no information loss has occurred **Note**: Information loss under these circumstances is a normal part of BGP4 and does not indicate an error.
Local Pref	The degree of preference for routes that use this set of attributes relative to other routes in the local AS.
Communities	The communities that routes with this set of attributes are in.
AS Path	The ASs through which routes with this set of attributes have passed. The local AS is shown in parentheses.

USING THE WEB MANAGEMENT INTERFACE

1. Select Show.
2. Select Attribute-Entries under BGP to display a panel such as the one shown in Figure 12.19. (The example shown here is for a routing switch that has been enabled for BGP4 but does not have any routes yet. During normal BGP4 operation, this panel shows numerous entries.)

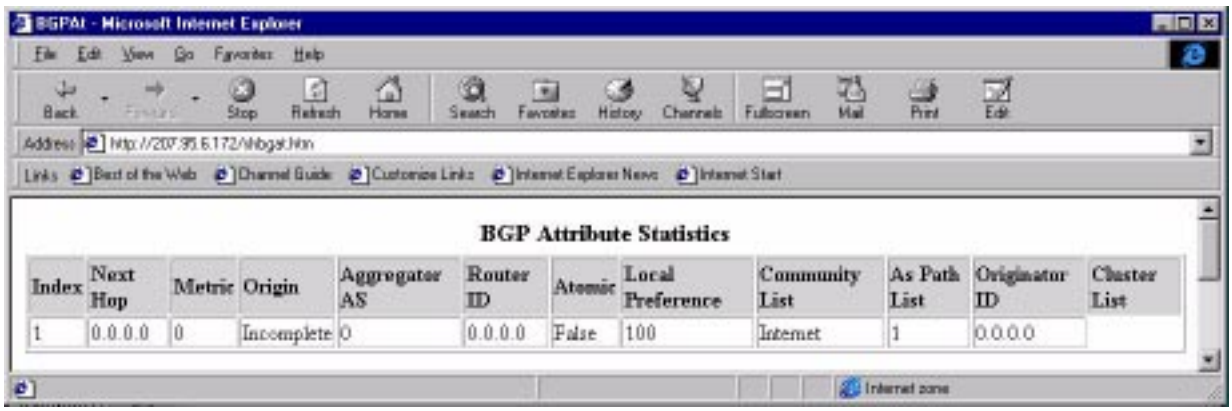

Figure 12.19 BGP Attribute Entries panel

Displaying the Routes BGP4 Has Placed in the IP Route Table

The IP route table indicates the routes it has received from BGP4 by listing “BGP” as the route type. You can view the IP route table using either of the following methods.

USING THE CLI

To display the IP route table, enter the following command:

```
HP9300# show ip route
```

Syntax: show ip route [<IP-addr> | <num> | bgp | ospf | rip]

Here is an example of the information displayed by this command. Notice that most of the routes in this example have type "B", indicating that their source is BGP4.

```
HP9300# show ip route
Total number of IP routes: 50834
B:BGP D:Directly-Connected  O:OSPF  R:RIP  S:Static
      Network Address  NetMask          Gateway          Port      Cost   Type
      3.0.0.0          255.0.0.0        192.168.13.2     1/1       0      B
      4.0.0.0          255.0.0.0        192.168.13.2     1/1       0      B
      9.20.0.0         255.255.128.0    192.168.13.2     1/1       0      B
      10.1.0.0         255.255.0.0      0.0.0.0          1/1       1      D
      10.10.11.0       255.255.255.0    0.0.0.0          2/24      1      D
      12.2.97.0        255.255.255.0    192.168.13.2     1/1       0      B
      12.3.63.0        255.255.255.0    192.168.13.2     1/1       0      B
      12.3.123.0       255.255.255.0    192.168.13.2     1/1       0      B
      12.5.252.0       255.255.254.0    192.168.13.2     1/1       0      B
      12.6.42.0        255.255.254.0    192.168.13.2     1/1       0      B
```

remaining 50824 entries not shown...

USING THE WEB MANAGEMENT INTERFACE

1. Select Show.
2. Select Routes under IP.

Clearing Traffic Counters

You can clear the counters (reset them to 0) for BGP4 messages. To do so, use one of the following methods.

USING THE CLI

To clear the BGP4 message counter for all neighbors, enter the following command:

```
HP9300# clear ip bgp traffic
```

syntax: clear ip bgp traffic

To clear the BGP4 message counter for a specific neighbor, enter a command such as the following:

```
HP9300# clear ip bgp neighbor 10.0.0.1 traffic
```

syntax: clear ip bgp neighbor <IP-addr> traffic

USING THE WEB MANAGEMENT INTERFACE

1. Select Clear to display the panel show in Figure 12.20.
2. Select one of the following options:
   - BGP Neighbor Traffic – clears the BGP4 message counters for all neighbors (the default) or a neighbor you select from the pulldown menu.
   - BGP Neighbor – clears the BGP4 message counters for all neighbors (the default) or a neighbor you select from the pulldown menu.
3. Select Apply.

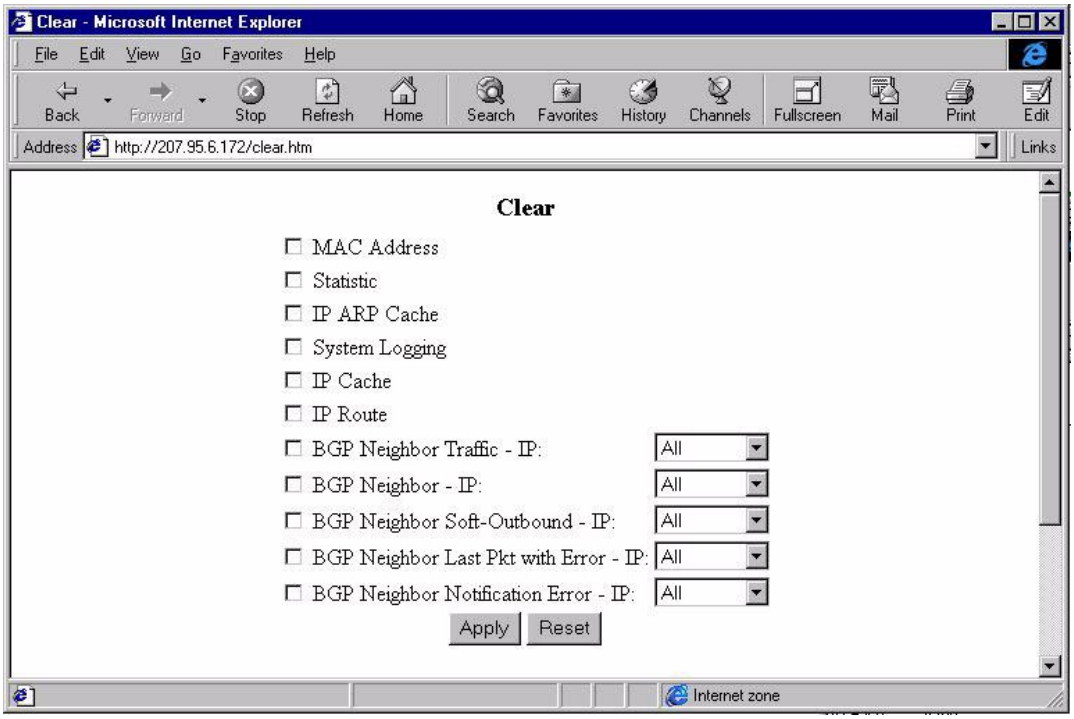

Figure 12.20 Web management Clear panel

Closing or Resetting Sessions With Neighbors

You can close a neighbor session or resend route updates to a neighbor. If you make changes to filters or route maps, use these methods to ensure that neighbors contain only the routes you want them to contain.

- If you close a neighbor session, the routing switch and the neighbor clear all the routes they learned from each other. When the routing switch and neighbor establish a new BGP4 session, they exchange route tables again. Use this method if you want the routing switch to relearn routes from the neighbor and resend its own route table to the neighbor.
- If you use the soft-outbound option, the routing switch compiles a list of all the routes it would normally send to the neighbor at the beginning of a session. However, before sending the updates, the 9304M, 9308M, and 6308M-SX routing switches also apply the filters and route maps you have configured to the list of routes. If the filters or route maps result in changes to the list of routes, the routing switch sends updates to advertise, change, or even withdraw routes on the neighbor as needed. This ensures that the neighbor receives only the routes you want it to contain. Even if the neighbor already contains a route learned from the routing switch that you later decided to filter out, using the soft-outbound option removes that route from the neighbor.

USING THE CLI

To close a neighbor session and thus flush all the routes exchanged by the routing switch and the neighbor, enter the following command:

```
HP9300# clear ip bgp neighbor all
```

To resend routes to a neighbor without closing the neighbor session, enter a command such as the following:

```
HP9300# clear ip bgp neighbor 10.0.0.1 soft-outbound
```

syntax: clear ip bgp neighbor all|<IP-addr> soft-outbound

USING THE WEB MANAGEMENT INTERFACE

To resend route information to a neighbor, use the following procedure:

1. Select Clear to display the panel show in Figure 12.20.
2. Select BGP Neighbor Soft-Outbound.
3. Use the default value All to resend the BGP4 route table to all neighbors or select a neighbor from the field's pulldown menu.
4. Select Apply.

Clearing Diagnostic Buffers

The routing switch stores the following BGP4 diagnostic information in buffers:

- The first 400 bytes of the last packet that contained an error
- The last NOTIFICATION message either sent or received by the routing switch

To display these buffers, use options with the **show ip bgp neighbors** command. See "Displaying BGP4 Neighbor Information" on page 12-46.

This information can be useful if you are working with HP Technical Support to resolve a problem. The buffers do not identify the system time when the data was written to the buffer. If you want to ensure that diagnostic data in a buffer is recent, you can clear the buffers. You can clear the buffers for a specific neighbor or for all neighbors.

If you clear the buffer containing the first 400 bytes of the last packet that contained errors, all the bytes are changed to zeros. The Last Connection Reset Reason field of the BGP neighbor table also is cleared.

If you clear the buffer containing the last NOTIFICATION message sent or received, the buffer contains no data.

USING THE CLI

To clear these buffers for neighbor 10.0.0.1, enter the following commands:

```
HP9300# clear ip bgp neighbor 10.0.0.1 last-packet-with-error
HP9300# clear ip bgp neighbor 10.0.0.1 notification-errors
```

syntax: clear ip bgp neighbor all|<IP-addr> <last-packet-with-error|notification-errors>

USING THE WEB MANAGEMENT INTERFACE

1. Select Clear to display the panel show in Figure 12.20.
2. Select one of the following:
   - BGP Neighbor Last Packet with Error – Clears the buffer containing the first 400 bytes of the last BGP4 packet that contained an error.
   - BGP Neighbor Notification Error – Clears the buffer containing the last NOTIFICATION message sent or received.
3. Select Apply.