Chapter 8
Configuring BGP4

This chapter provides details on how to configure ***Border Gateway Protocol version 4 (BGP4)*** on the HP 9304M, HP 9308M, and HP 6208M-SX routing switches using the CLI and the Web management interface.

BGP4 is described in RFC 1771. The BGP4 implementation on the HP 9304M, HP 9308M, and HP 6208M-SX routing switches fully complies with RFC 1771. The implementation also includes RFC 1745 (OSPF Interactions), RFC 1965 (BGP4 Confederations), and RFC 1997 (BGP Communities Attributes).

To display BGP4 configuration information and statistics, see "Displaying BGP4 Information" on page 8-76.

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 the *Command Line Interface Reference*.

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

NOTE: The HP 6308M-SX and Chassis devices that use non-redundant 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 "Displaying and Modifying System Parameter Default Settings" in the "Configuring Basic Features" chapter of Book 1.

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 HP 9304M, HP 9308M, and HP 6208M-SX routing switches.

Figure 8.1 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 HP 9304M, HP 9308M, and HP 6208M-SX routing switches can be configured to redistribute routes among BGP4, RIP, and OSPF. They also can redistribute static routes.

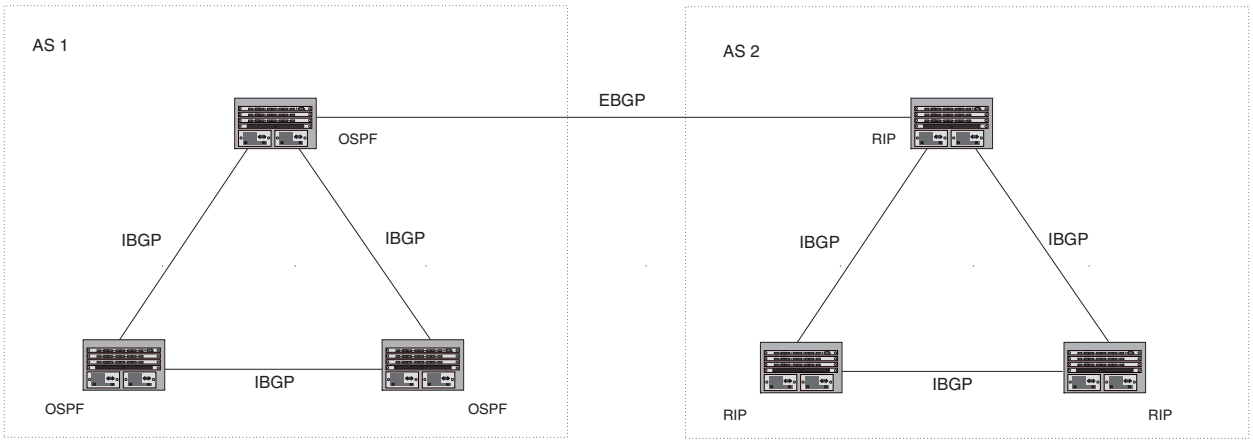

Figure 8.1 Example BGP4 ASs

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 8-4 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 8-20.)

- 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”.
- Confederation membership.
- Closest IBGP neighbor – The closest internal path to the destination within the local AS.
- Number of paths available for load sharing.

The HP 9304M, HP 9308M, and HP 6208M-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. All paths within a confederation have the same length.
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 from a BGP neighbor outside of the confederation
   - Routes received through EBGP from a BGP router within the confederation
   - 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.

NOTE: HP routing switches support BGP4 load sharing among up to eight equal paths. BGP4 load sharing enables the routing switch to balance the traffic across the multiple paths instead of choosing just one path based on router ID. See “Changing the Maximum Number of Paths for BGP4 Load Sharing” on page 8-24 for more information.

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 HP 9304M, HP 9308M, and HP 6208M-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, you can configure the router 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 disabled by default.
- BGP Identifier – The router ID. The BGP Identifier (router ID) identifies the BGP4 router to other BGP4 routers. HP routing switches use the same router ID for OSPF and BGP4. If you do not set a router ID, the software uses the IP address on the lowest numbered loopback interface configured on the router. If the routing switch does not have a loopback interface, the default router ID is the lowest numbered IP address configured on the device. For more information or to change the router ID, see “Changing the Router ID” on page 5-8.
- 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 an HP 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.

   NOTE: You must specify the local AS number. The routing switch does not enable BGP4 until you specify the local AS number.

3. Add each BGP4 neighbor (peer BGP4 router) and identify the AS the neighbor is in.
4. Save the BGP4 configuration information to the system configuration file.

NOTE: By default, the HP router ID is the IP address configured on the lowest numbered loopback interface. If the routing switch does not have a loopback interface, the default router ID is the lowest numbered IP address configured on the device. For more information or to change the router ID, see "Changing the Router ID" on page 5-8.

USING THE CLI

```
HP9300> enable
HP9300# configure terminal
HP9300(config)# router bgp
BGP4: Please configure 'local-as' parameter in order to enable BGP4.
HP9300(config-bgp-router)# local-as 10
HP9300(config-bgp-router)# neighbor 209.157.23.99 remote-as 100
HP9300(config-bgp-router)# write memory
```

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the Enable radio button next to BGP.
3. Enter the local AS number in the Local AS field.
4. Click the Apply button to apply the changes to the device's running-config file.
5. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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.)
- Required – Specify the local AS number.
- 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 – Enable 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 – Enable 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 – Configure the routing switch as a member of a BGP4 confederation.
- 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 – Change the number of paths for BGP4 load sharing.
- 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.
- Optional – Define route flap dampening parameters.

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 BGP configuration panel, shown in Figure 8.2. You can access all other parameters using links on the BGP configuration panel or from the Configure->BGP options in the tree view. Select Configure->BGP-General to display the BGP configuration panel.

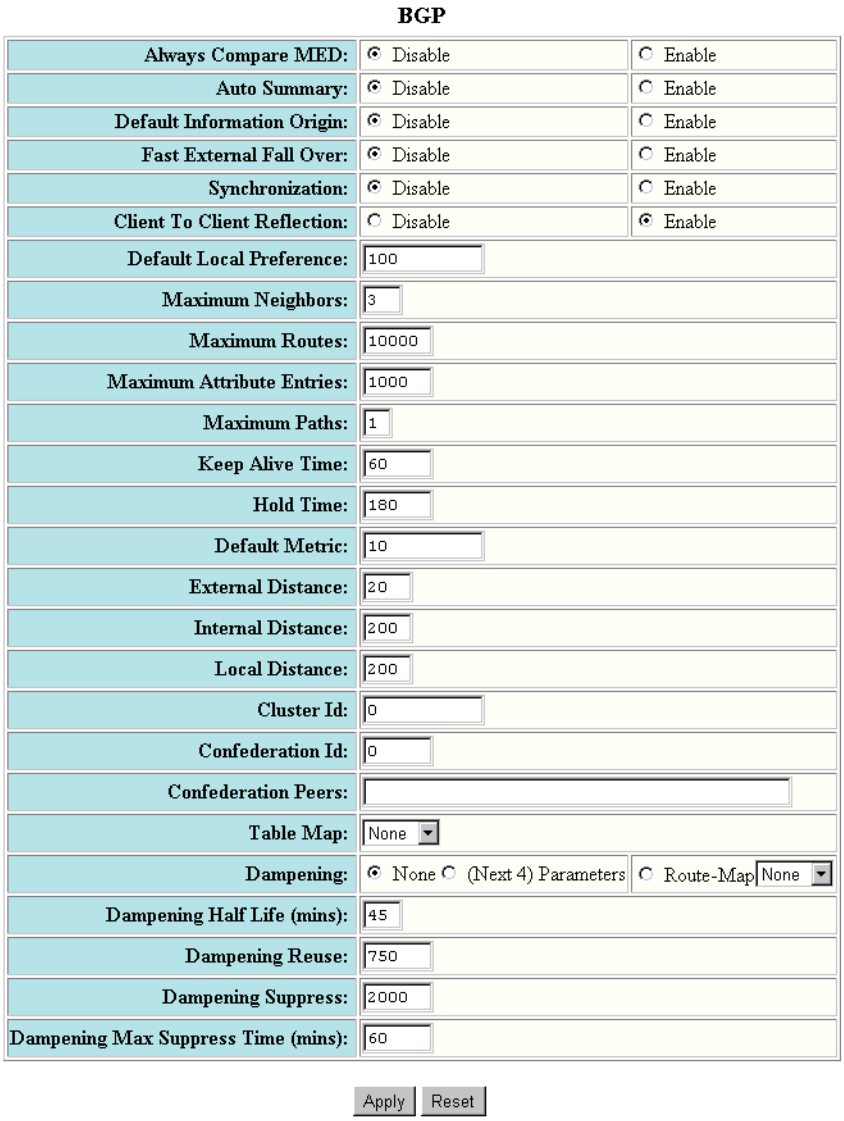

Figure 8.2 BGP configuration panel

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.
- Set or change the local AS.
- Add neighbors.
- Disable or enable 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.
- Disable or enable IGP and BGP4 synchronization.
- Enable or disable auto summary.
- Change the default metric.
- Disable or re-enable route reflection.
- Configure confederation parameters.
- Disable or re-enable load sharing.
- Change the maximum number of load-sharing paths.
- Define route flap dampening parameters.

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 8-95.

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 following table lists the new default maximum number of neighbors, routes, and route attributes on HP 9304M or HP 9308M routing switches using the standard management module, 128 MB redundant management modules, and 256 MB redundant management modules.

Platform	Default Maximum BGP4 Neighbors	Default Maximum BGP4 Routes	Default Maximum BGP4 Route Attributes
Standard management module with 32 MB	3	3	3
Redundant management module with 128 MB	15	200,000	100,000
Redundant management module with 256 MB	50	250,000	150,000

The following table lists the default maximum number of neighbors, routes, and route attributes you can configure on HP 9304M or HP 9308M routing switches. To use maximum values that are higher than the defaults, you must reconfigure the device's memory for the higher amount.

Platform	Maximum BGP4 Neighbors	Maximum BGP4 Routes	Maximum BGP4 Route Attributes
Standard (non-redundant) management module with 32 MB	3	30,000	10,000
Redundant management module with 128 MB	15	500,000	200,000
Redundant management module with 256 MB	50	1,000,000	200,000

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.

See "Changing the Maximum Number of Neighbors" on page 8-21 and "Changing the Maximum Number of Routes" on page 8-22 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. Set the local AS number.
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 enable fast external fallover.
7. Optionally change the maximum number of BGP4 neighbors, routes, or route-attribute entries the routing switch can have.

8. Optionally change the maximum number of BGP4 load sharing paths.
9. Optionally specify a list of individual networks in the local AS to be advertised to remote ASs using BGP4.
10. Optionally change the default local preference, default information originate, default MED (metric), or administrative distances. (You change these parameters independently of one another.)
11. Optionally configure the routing switch to always compare MEDs (metrics) when choosing a route.
12. Optionally enable synchronization of routes between BGP4 and IGP.
13. Optionally enable automatic summarization of subnets at the classical IP boundaries (classes A, B, and C).
14. Optionally aggregate routes in the BGP4 route table into CIDR blocks.
15. Optionally configure the routing switch as a BGP4 route reflector.
16. Optionally configure the routing switch as a member of a BGP4 confederation.
17. Optionally change the default metric for routes that BGP4 redistributes into RIP or OSPF.
18. Optionally define BGP4 address filters, AS-path filters, or community filters.
19. Optionally define BGP4 route map entries.
20. Optionally define route flap dampening parameters.
21. 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 8-20.

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> enable
HP9300# configure terminal
HP9300(config)# router bgp
BGP4: Please configure 'local-as' parameter in order to enable BGP4.
HP9300(config-bgp-router)# local-as 10
HP9300(config-bgp-router)# neighbor 209.157.23.99 remote-as 100
HP9300(config-bgp-router)# write memory
```

Syntax: [no] local-as <num>

The <num> parameter specifies the local AS number.

These commands place you in the BGP4 router level and allow you to configure the local AS number, which is required to enable BGP4 and modify parameters.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the Enable radio button next to BGP.
3. Enter the local AS number in the Local AS field.
4. Click the Apply button to apply the changes to the device’s running-config file.
5. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

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 HP 9304M, HP 9308M, and HP 6208M-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 change 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 HP routing switch, but do not specify an IP address in use by another device.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to IP in the tree view to expand the list of IP option links.
4. Click on the General link to display the IP configuration panel.
5. Edit the value in the Router ID field. Specify a valid IP address that is not in use on another device in the network.
6. Click the Apply button to save the change to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 and are not advertised to the Internet community.

To set the local AS number, use either of the following methods.

USING THE CLI

To set the local AS number, enter commands such as the following:

```
HP9300(config)# router bgp
BGP4: Please configure 'local-as' parameter in order to enable BGP4.
HP9300(config-bgp-router)# local-as 10
HP9300(config-bgp-router)# write memory
```

Syntax: [no] local-as <num>

The <num> parameter specifies the local AS number.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the Enable radio button next to BGP.
3. Enter the local AS number in the Local AS field.
4. Click the Apply button to apply the changes to the device's running-config file.
5. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Adding a Loopback Interface

You can configure the routing switch to use a loopback interface instead of a specific port to communicate with a BGP4 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 BGP4 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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the IP Address link to display a table listing the configured IP addresses.
3. Select the Loop Back link.

   NOTE: If the device already has loopback interfaces, a table listing the interfaces is displayed. Click the Modify button to the right of the row describing an interface to change its configuration, or click the Add Loop Back link to display the Router Loop Back configuration panel.

   ---

4. Select the loopback interface number from the Loopback field's pulldown menu. You can select from 1 – 8.
5. Select the status. The interface is enabled by default.
6. Click Add to add the new interface.
7. Click on Configure in the tree view to display the configuration options.

8. Click on IP to display the IP configuration options.
9. Select the Add IP Address link to display the Router IP Address panel.
10. Select the loopback interface from the Port field's pulldown menu. For example, to select loopback interface 1, select "lb1". (If you are configuring a Chassis device, you can have any slot number in the Slot field. Loopback interfaces are not associated with particular slots or physical ports.)
11. Enter the loopback interface's IP address in the IP Address field.
12. Enter the network mask in the Subnet Mask field.
13. Click the Add button to save the change to the device's running-config file.
14. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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.

NOTE: The routing switch attempts to establish a BGP4 session with a neighbor as soon as you enter a command specifying the neighbor's IP address. If you want to completely configure the neighbor parameters before the routing switch establishes a session with the neighbor, you can administratively shut down the neighbor. See "Administratively Shutting Down a Session with a BGP4 Neighbor" on page 8-19.

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 <ip-addr> remote-as <as-number>
[advertisement-interval <num>] [default-originate [route-map <map-name>]]
[description <string>] [distribute-list in | out <num,num,...>] [ebgp-multihop [<num>]]
[filter-list in | out | weight <num,num,...>] [maximum-prefix <num>] [next-hop-self]
[password <string>] [prefix-list <string>] [remote-as <as-number>] [remove-private-as]
[route-map in | out <map-name>] [route-reflector-client] [send-community] [shutdown]
[timers keep-alive <num> hold-time <num>] [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 a 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.

default-originate [**route-map** <map-name>] configures the **routing switch** to send the default route 0.0.0.0 to the neighbor. If you use the **route-map** <map-name> parameter, the route map injects the default route conditionally, based on the match conditions in the route map.

description <string> specifies a name for the neighbor. You can enter an alphanumeric text string up to 80 characters long.

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 router 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 8-40.

ebgp-multihop [<num>] 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. The <num> parameter specifies the TTL you are adding for the neighbor. You can specify a number from 0 – 255. The default is 0. If you leave the EBGP TTL value set to 0, the software uses the IP TTL value.

filter-list in | out | weight <num,num,...> specifies an AS-path filter list or a list of AS-path Access Control Lists (ACLs). The **in | out** keyword specifies whether the list is applied on updates received from the neighbor or sent to the neighbor. If you specify **in** or **out**, The <num,num,...> parameter specifies the list of AS-path filters. The router 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. The **weight** <num> parameter specifies a weight that the routing switch applies to routes received from the neighbor that match the AS-path filter or ACL. You can specify a number from 0 – 65535.

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

NOTE: The AS-path filter or ACL must already be configured. See "Filtering AS-Paths" on page 8-41.

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 HP 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 8-22.

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

password <string> specifies an MD5 password for securing sessions between the **routing switch** and the neighbor. You can enter a string up to 80 characters long. The string can contain any alphanumeric characters, but the first character cannot be a number. If the password contains a number, do not enter a space following the number.

prefix-list <string> specifies an IP prefix list. You can use IP prefix lists to control routes to and from the neighbor. IP prefix lists are an alternative method to AS-path filters. You can configure up to 1000 prefix list filters. The filters can use the same prefix list or different prefix lists. To configure an IP prefix list, see "Defining IP Prefix Lists" on page 8-50.

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 router to remove private AS numbers from UPDATE messages the router sends to this neighbor. The router 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 8-55.

route-reflector-client specifies that this neighbor is a route-reflector client of the router. Use the parameter only if this router is going to be a route reflector. For information, see "Configuring Route Reflection Parameters" on page 8-31. This option is disabled by default.

send-community enables sending the community attribute in updates to the specified neighbor. By default, the router does not send the community attribute.

shutdown administratively shuts down the session with this neighbor. Shutting down the session allows you to completely configure the neighbor and save the configuration without actually establishing a session with the neighbor. This option is disabled by default.

timers keep-alive <num> **hold-time** <num> overrides the global settings for the Keep Alive Time and Hold Time. For the Keep Alive Time, you can specify from 0 – 65535 seconds. For the Hold Time, you can specify 0 or 3 – 65535 (1 and 2 are not allowed). If you set the Hold Time to 0, the router waits indefinitely for messages from a neighbor without concluding that the neighbor is dead. The defaults for these parameters are the currently configured global Keep Alive Time and Hold Time. For more information about these parameters, see "Changing the Keep Alive Time and Hold Time" on page 8-20.

update-source loopback <num> configures the router 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 router 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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Neighbor link to display the BGP Neighbor panel.

NOTE: If the device already has neighbors, a table listing the neighbors is displayed. Click the Modify button to the right of the row describing the neighbor to change its configuration, or click the Add Neighbor link to display the BGP Neighbor configuration panel.

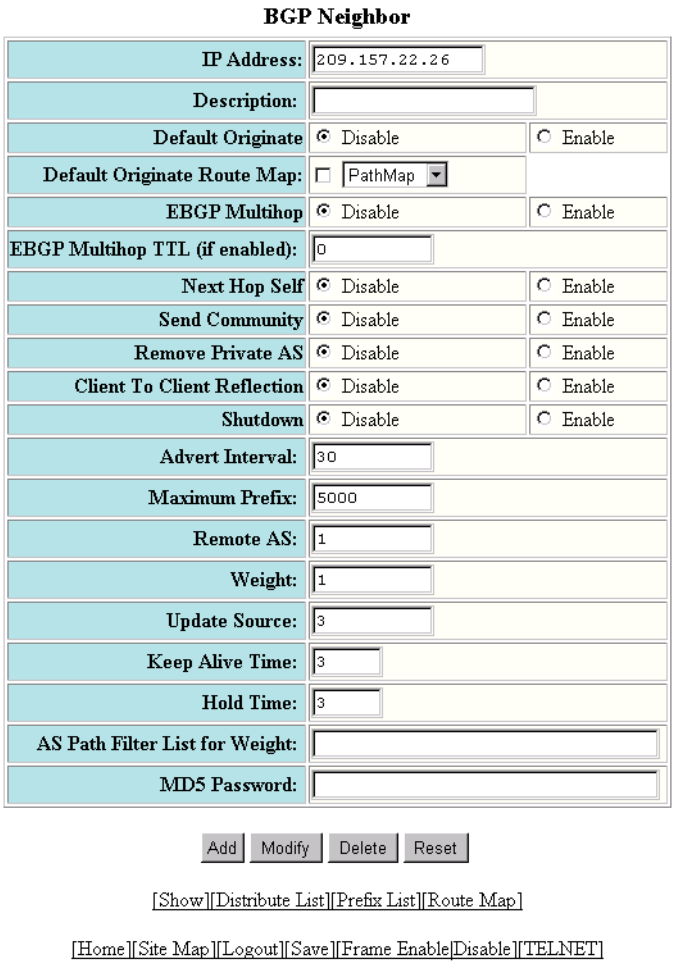

1. Enter the neighbor's IP address in the IP Address field.
2. Enter a description in the Description field.
3. Select Enable next to Default Originate if you want to enable this feature for the neighbor. By default, the routing switch does not advertise a default route using BGP4. A BGP4 default route 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. Select the checkbox next to Default Originate Route Map and select a route map from the pulldown menu if you want to use a route map to control advertisement of default routes.
5. Select Enable next to EBGP Multihop if the neighbor is multiple EBGP hops away.
6. If you enabled EBGP Multihop, enter the TTL for EBGP multihop in the EBGP Multihop TTL field. You can specify a number from 0 – 255. The default is 0. If you leave the EBGP TTL value set to 0, the software uses the IP TTL value.
7. Select Enable next to Next Hop Self if the router should list itself as the next hop in updates sent to the neighbor. This option is disabled by default.
8. Select Enable next to Send Community if you want to send the community attribute in updates to the neighbor. By default, the router does not send the community attribute.
9. Select Enable next to Remove Private AS if you want the router to remove private AS numbers from UPDATE messages the router sends to this neighbor. The router 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.

10. Select Enable next to Client To Client Reflection if this neighbor is a route-reflector client of the router. Use the parameter only if this router is going to be a route reflector. For information, see “Configuring Route Reflection Parameters” on page 8-31. This option is disabled by default.
11. Select Enable next to Shutdown if you want to administratively shut down the session with this neighbor. Shutting down the session allows you to completely configure the neighbor and save the configuration without actually establishing a session with the neighbor. This option is disabled by default.
12. Enter the advertisement interval in the Advert Interval field. This parameter 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.
13. Edit the value in the Maximum Prefix field to change the maximum prefix. The maximum prefix is 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 HP 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 8-22.
14. Enter the remote AS number in the Remote AS field. The remote AS number is the number of the AS the neighbor is in.
15. Enter the weight you want the routing switch to add to routes received from the specified neighbor. BGP4 prefers larger weights over smaller weights. The default weight is 0.
16. Enter the number of an update source loopback interface in the Update Source field. This parameter configures the router 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 router 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 – 8. There is no default.
17. Enter a Keep Alive time in the Keep Alive Time field. This parameter overrides the global BGP4 Keep Alive Time configured on the routing switch. You can specify from 0 – 65535 seconds. The default is the current global setting.
18. Enter a Hold Time in the Hold Time field. This parameter overrides the global BGP4 Hold Time configured on the routing switch. You can specify 0 or 3 – 65535 (1 and 2 are not allowed). If you set the Hold Time to 0, the router waits indefinitely for messages from a neighbor without concluding that the neighbor is dead. The default is the current global setting.

    NOTE: Set the Hold Time to three times the value of the Keep Alive Time. For information about these parameters, see “Changing the Keep Alive Time and Hold Time” on page 8-20.

19. If you specified a weight in the Weight field, enter a list of AS Path filters in the AS Path Filter List for Weight field. The router 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 8-41.

20. Enter a password in the MD5 Password field to secure the routing switch’s sessions with this neighbor.

    NOTE: You must configure the neighbor to use the same password.

21. Click the Add button (if you are adding a new neighbor) or the Modify button (if you are modifying a neighbor that is already configured) to apply the changes to the device’s running-config file.

22. Select the Save link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Administratively Shutting Down a Session with a BGP4 Neighbor

You can prevent the routing switch from starting a BGP4 session with a neighbor by administratively shutting down the neighbor. This option is very useful for situations in which you want to configure parameters for a neighbor but are not ready to use the neighbor. You can shut the neighbor down as soon as you have added it the routing switch, configure the neighbor parameters, then allow the routing switch to reestablish a session with the neighbor by removing the shutdown option from the neighbor.

When you apply the new option to shut down a neighbor, the option takes place immediately and remains in effect until you remove the option. If you save the configuration to the startup-config file, the shutdown option remains in effect even after a software reload.

NOTE: The software also contains an option to end the session with a BGP4 neighbor and thus clear the routes learned from the neighbor. Unlike this clear option, the option for shutting down the neighbor can be saved in the startup-config file and thus can prevent the routing switch from establishing a BGP4 session with the neighbor even after reloading the software.

NOTE: If you notice that a particular BGP4 neighbor never establishes a session with the HP routing switch, check the routing switch's running-config and startup-config files to see whether the configuration contains a command that is shutting down the neighbor. The neighbor may have been shut down previously by an administrator.

To shut down a BGP4 neighbor, use either of the following methods.

USING THE CLI

To shut down a BGP4 neighbor, enter commands such as the following:

```
HP9300(config)# router bgp
HP9300(config-bgp-router)# neighbor 209.157.22.26 shutdown
HP9300(config-bgp-router)# write memory
```

Syntax: [no] neighbor <ip-addr> shutdown

The <ip-addr> parameter specifies the IP address of the neighbor.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Neighbor link to display the BGP Neighbor panel.

   NOTE: If the device already has neighbors, a table listing the neighbors is displayed. Click the Modify button to the right of the row describing the neighbor to change its configuration, or click the Add Neighbor link to display the BGP Neighbor configuration panel.

   ---

5. Enter or modify parameters as needed. For detailed information, see "Adding BGP4 Neighbors" on page 8-14.
6. Select the Enable radio button next to Shutdown.
7. Click the Add button (if you are adding a new neighbor) or the Modify button (if you are modifying a neighbor that is already configured) to apply the changes to the device's running-config file.
8. Select the Save link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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.

NOTE: You can override the global Keep Alive Time and Hold Time on individual neighbors. See “Adding BGP4 Neighbors” on page 8-14.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Edit the number in the Keep Alive Time field. The Keep Alive Time can be 0 – 65535.
6. Edit the number in the Hold Time field. The Hold Time can be 0 or 3 – 65535 (1 and 2 are not allowed). If you set the Hold Time to 0, the router waits indefinitely for messages from a neighbor without concluding that the neighbor is dead.

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

7. Click the Apply button to apply the changes to the device’s running-config file.
8. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Enabling 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 router 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.

The router waits for the Hold Time to expire before ending the connection to a directly-attached BGP4 neighbor that dies.

For directly attached neighbors, the router to immediately senses loss of a connection to the neighbor from a change to the state of the port or interface that connects the router to its neighbor. For directly attached EBGP neighbors, the router 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 enable the router to immediately close the BGP4 session and TCP connection to locally attached neighbors that die, use either of the following methods.

USING THE CLI

To enable fast external fallover, enter the following command:

```
HP9300(config-bgp-router)# fast-external-fallover
```

To disable fast external fallover again, enter the following command:

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

Syntax: [no] fast-external-fallover

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Select Disable or Enable next to Fast External Fall Over.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Changing the Maximum Number of Neighbors

You can change the maximum number of BGP4 neighbors the routing switch can have 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 8-31.

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. See "Memory Considerations" on page 8-10 for the maximum for your device. The change takes effect after the router is rebooted.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Change the number in the Maximum Neighbors field. The maximum number you can enter depends on the device you are configuring.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
8. Click on Command in the tree view to list the command options.
9. Select the Reload link and click on Yes when prompted. You must reload the software to place this configuration change into effect.

Changing the Maximum Number of Routes

You can change the maximum number of BGP4 routes the routing switch can have using either of the following methods.

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 8-14.

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. See "Memory Considerations" on page 8-10 for the maximum for your device. The change takes effect after the router is rebooted.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Change the number in the Maximum Routes field. The maximum number you can enter depends on the device you are configuring.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
8. Click on Command in the tree view to list the command options.
9. Select the Reload link and click on Yes when prompted. You must reload the software to place this configuration 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 8-90 for a description of the route-attribute fields shown in the example above.

You can change the maximum number of route-attribute entries the routing switch can contain using either of the following methods.

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 router. See “Memory Considerations” on page 8-10 for the maximum for your device. The change takes effect after the router is rebooted.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Change the number in the Maximum Attribute Entries field. The maximum number you can enter depends on the device you are configuring.
6. Click the Apply button to apply the changes to the device’s running-config file.

7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
8. Click on Command in the tree view to list the command options.
9. Select the Reload link and click on Yes when prompted. You must reload the software to place this configuration change into effect.

Changing the Maximum Number of Paths for BGP4 Load Sharing

Load sharing enables the routing switch to balance traffic to a route across multiple equal-cost paths of the same type (EBGP or IBGP) for the route.

To configure the routing switch to perform BGP4 load sharing:

- Enable IP load sharing if it is disabled.
- Set the maximum number of paths. The default maximum number of BGP4 load sharing paths is 1, which means no BGP4 load sharing takes place by default.

NOTE: The maximum number of BGP4 load sharing paths cannot be greater than the maximum number of IP load sharing paths.

How Load Sharing Affects Route Selection

During evaluation of multiple paths to select the best path to a given destination for installment in the IP route table, the last comparison the routing switch performs is a comparison of the internal paths.

- When IP load sharing is disabled, the routing switch prefers the path to the router with the lower router ID.
- When IP load sharing is enabled, the routing switch balances the traffic across the multiple paths instead of choosing just one path based on router ID.

See "How BGP4 Selects a Path for a Route" on page 8-3 for a description of the BGP4 algorithm.

When you enable IP load sharing, the routing switch can load balance BGP4 or OSPF routes across up to four equal paths by default. You can change the number of IP load sharing paths to a value from 2 – 8.

How Load Sharing Works

Load sharing is performed in round-robin fashion and is based on the destination IP address only. The first time the router receives a packet destined for a specific IP address, the router uses a round-robin algorithm to select the path that was not used for the last newly learned destination IP address. Once the router associates a path with a particular destination IP address, the router will always use that path as long as the router contains the destination IP address in its cache.

NOTE: The routing switch does not perform source routing. The router is concerned only with the paths to the next-hop routers, not the entire paths to the destination hosts.

Changing the Maximum Number of Shared BGP4 Paths

When IP load sharing is enabled, BGP4 can balance traffic to a specific destination across up to four equal paths. You can set the maximum number of paths to a value from 1 – 4. The default is 1.

NOTE: The maximum number of BGP4 load sharing paths cannot be greater than the maximum number of IP load sharing paths. To increase the maximum number of IP load sharing paths, use the **ip load sharing** <num> command at the global CONFIG level of the CLI or use the # of Paths field next to Load Sharing on the IP configuration panel of the Web management interface.

USING THE CLI

To change the maximum number of shared paths, enter commands such as the following:

```
HP9300(config)# router bgp
HP9300(config-bgp-router)# maximum-paths 4
HP9300(config-bgp-router)# write memory
```

Syntax: [no] maximum-paths <num>

The <num> parameter specifies the maximum number of paths across which the routing switch can balance traffic to a given BGP4 destination. You can change the maximum number of paths to a value from 2 – 4. The default is 1.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Edit the number in the # of Paths field if needed. You can specify from 1 – 8 paths. The default is 1. You cannot set the maximum number of BGP4 paths to a number higher than the IP load sharing maximum number of paths.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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. You can specify up to 600 networks.

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 <ip-mask> [weight <num>] [backdoor]

The <ip-addr> is the network number and the **mask** <ip-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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.

4. Click on the Network link.
   - If the device does not have any BGP networks configured, the BGP Network configuration panel is displayed, as shown in the following example.
   - If a BGP network is already configured and you are adding a new one, click on the Add Network link to display the BGP Network configuration panel, as shown in the following example.
   - If you are modifying an existing BGP network, click on the Modify button to the right of the row describing the network to display the BGP Network configuration panel, as shown in the following example.

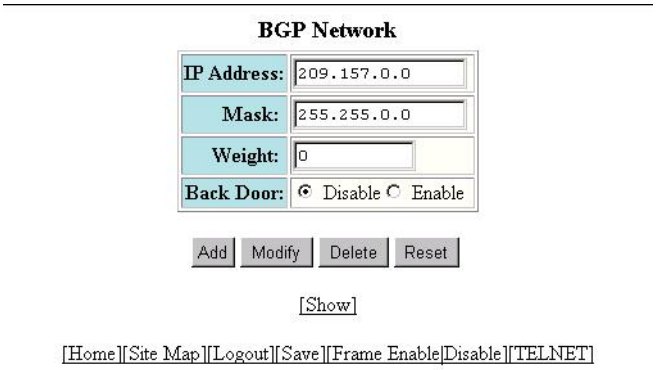

5. Enter the network address in the IP Address field.
6. Enter the network mask in the Mask field.
7. Optionally enter a weight to be added to routes to this network.
8. If you want to tag the route as a backdoor route, select Enable next to Back Door.
9. Click the Apply button to apply the changes to the device's running-config file.
10. Select the Save link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 8-55. See "How BGP4 Selects a Path for a Route" on page 8-3 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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Change the number in the Default Local Preference field. You can enter a number from 0 – 4294967295.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Select Disable or Enable next to Default Information Originate.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Change the number in the Default Metric field. You can enter a number from 0 – 4294967295.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 HP 9304M, HP 9308M, and HP 6208M-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 6-36.
- To change the default administrative distance for RIP, see "Modifying the Default Administrative Distance" on page 5-40.
- To change the default administrative distance for static routes, see "Defining Static IP Routes" on page 5-14.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Change the number in the External Distance field to change the EBGP distance. You can enter a number from 1 – 255.
6. Change the number in the Internal Distance field to change the IBGP distance. You can enter a number from 1 – 255.
7. Change the number in the Local Distance field to change the local distance. You can enter a number from 1 – 255.
8. Click the Apply button to apply the changes to the device's running-config file.
9. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Select Disable or Enable next to Always Compare MED.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Synchronizing Routes

By default, the routing switch does not wait until the IGPs in the local AS have fully exchanged route information before BGP4 advertises the routes to its remote BGP4 neighbors. The routing switch advertises routes to its remote BGP4 neighbors regardless of whether the routes are learned or have already been propagated throughout the local AS.

If you want the routing switch to wait until the IGPs in the local AS have fully exchanged route information before BGP4 advertises the routes to its remote BGP4 neighbors, enable synchronization.

To enable synchronization, use either of the following methods.

USING THE CLI

To enable synchronization, enter the following command:

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

To disable synchronization again, enter the following command:

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

Syntax: [no] synchronization

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Select Disable or Enable next to Synchronization.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 disabled by default. If you want to enable 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 8-35.

USING THE CLI

To enable auto summary, enter the following command:

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

To disable auto summary again, enter the following command:

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

Syntax: [no] auto-summary

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Select Disable or Enable next to Auto Summary.
6. Click the Apply button to apply the changes to the device's running-config file.
7. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 many 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 HP 9304M, HP 9308M, and HP 6208M-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 8.3 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.

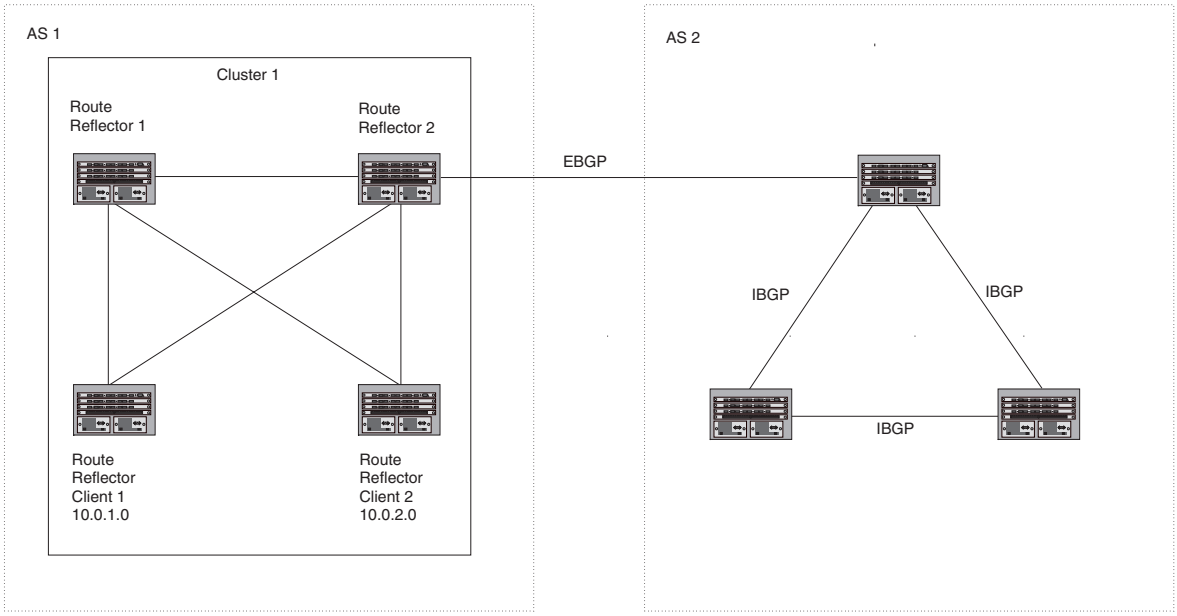

Figure 8.3 Example route reflector configuration

To configure an HP 9304M, HP 9308M, and HP 6208M-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 8.3. 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 8-14.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to System in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. If route reflection is not already enabled, select Enable next to Client To Client Reflection.
6. 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.
7. Click the Apply button to apply the changes to the device’s running-config file.
8. Click on the Neighbor link at the bottom of the BGP configuration panel or under BGP in the Configure section of the tree view.
9. 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. The BGP configuration panel is displayed.
10. 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. See “Adding BGP4 Neighbors” on page 8-14 for information about the other neighbor parameters.
11. Click the Add button to apply the changes to the device’s running-config file.
12. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Configuring Confederations

A ***confederation*** is a BGP4 Autonomous System (AS) that has been subdivided into multiple, smaller ASs. Subdividing an AS into smaller ASs simplifies administration and reduces BGP-related traffic, thus reducing the complexity of the Interior Border Gateway Protocol (IBGP) mesh among the BGP routers in the AS.

The HP implementation of this feature is based on RFC 1965.

Normally, all BGP routers within an AS must be fully meshed, so that each BGP router has interfaces to all the other BGP routers within the AS. This is feasible in smaller ASs but becomes unmanageable in ASs containing many BGP routers.

When you configure BGP routers into a confederation, all the routers within a sub-AS (a subdivision of the AS) use IBGP and must be fully meshed. However, routers use EBGP to communicate between different sub-ASs.

NOTE: Another method for reducing the complexity of an IBGP mesh is to use route reflection. However, if you want to run different Interior Gateway Protocols (IGPs) within an AS, configure a confederation. You can run a separate IGP within each sub-AS.

To configure a confederation, configure groups of BGP routers into sub-ASs. A sub-AS is simply an AS. The term "sub-AS" distinguishes ASs within a confederation from ASs that are not in a confederation. For the viewpoint of remote ASs, the confederation ID is the AS ID. Remote ASs do not know that the AS represents multiple sub-ASs with unique AS IDs.

NOTE: You can use any valid AS numbers for the sub-ASs. If your AS is connected to the Internet, HP recommends that you use numbers from within the private AS range (64512 – 65535). These are private ASs numbers and BGP4 routers do not propagate these AS numbers to the Internet.

Perform the following configuration tasks on each BGP router within the confederation:

- Configure the local AS number. The local AS number indicates membership in a sub-AS. All BGP routers with the same local AS number are members of the same sub-AS. BGP routers use the local AS number when communicating with other BGP routers within the confederation.
- Configure the confederation ID. The confederation ID is the AS number by which BGP routers outside the confederation know the confederation. Thus, a BGP router outside the confederation is not aware and does not care that your BGP routers are in multiple sub-ASs. BGP routers use the confederation ID when communicating with routers outside the confederation. The confederation ID must be different from the sub-AS numbers.
- Configure the list of the sub-AS numbers that are members of the confederation. All the routers within the same sub-AS use IBGP to exchange router information. Routers in different sub-ASs within the confederation use EBGP to exchange router information.

Configuring a BGP Confederation

To configure a routing switch to be a member of a BGP confederation, use one of the following methods.

USING THE CLI

To configure four routing switches to be a member of confederation 10, consisting of two sub-ASs (64512 and 64513), enter commands such as the following.

Commands for Router A

```
HP9300A(config)# router bgp
HP9300A(config-bgp-router)# local-as 64512
HP9300A(config-bgp-router)# confederation identifier 10
HP9300A(config-bgp-router)# confederation peers 64512 64513
HP9300A(config-bgp-router)# write memory
```

Syntax: local-as <num>

The <num> parameter with the **local-as** command indicates the AS number for the BGP routers within the sub-AS. You can specify a number from 1 – 65535. HP recommends that you use a number within the range of well-known private ASs, 64512 – 65535.

Syntax: confederation identifier <num>

The <num> parameter with the **confederation identifier** command indicates the confederation number. The confederation ID is the AS number by which BGP routers outside the confederation know the confederation. Thus, a BGP router outside the confederation is not aware and does not care that your BGP routers are in multiple sub-ASs. BGP routers use the confederation ID when communicating with routers outside the confederation. The confederation ID must be different from the sub-AS numbers. You can specify a number from 1 – 65535.

Syntax: confederation peers <num> [<num> …]

The <num> parameter with the **confederation peers** command indicates the sub-AS numbers for the sub-ASs in the confederation. You must specify all the sub-ASs contained in the confederation. All the routers within the same sub-AS use IBGP to exchange router information. Routers in different sub-ASs within the confederation use EBGP to exchange router information. You can specify a number from 1 – 65535.

Commands for Router B

```
HP9300B(config)# router bgp
HP9300B(config-bgp-router)# local-as 64512
HP9300B(config-bgp-router)# confederation identifier 10
HP9300B(config-bgp-router)# confederation peers 64512 64513
HP9300B(config-bgp-router)# write memory
```

Commands for Router C

```
HP9300C(config)# router bgp
HP9300C(config-bgp-router)# local-as 64513
HP9300C(config-bgp-router)# confederation identifier 10
HP9300C(config-bgp-router)# confederation peers 64512 64513
HP9300C(config-bgp-router)# write memory
```

Commands for Router D

```
HP9300D(config)# router bgp
HP9300D(config-bgp-router)# local-as 64513
HP9300D(config-bgp-router)# confederation identifier 10
HP9300D(config-bgp-router)# confederation peers 64512 64513
HP9300D(config-bgp-router)# write memory
```

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel, shown in Figure 8.2 on 8-8.
5. Enter the confederation ID in the Confederation ID field. The confederation ID must be different from the sub-AS numbers. You can specify a number from 1 – 65535.
6. Enter the AS numbers of the peers (sub-ASs) within the confederation in the Confederation Peers field. Separate the AS numbers with spaces. You must specify all the sub-ASs contained in the confederation. All the routers within the same sub-AS use IBGP to exchange router information. Routers in different sub-ASs within the confederation use EBGP to exchange router information. You can specify a number from 1 – 65535.
7. Click the Apply button to apply the changes to the device's running-config file.
8. Select the Save link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 8-55 for information on defining a route map.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Aggregate Address link to display the BGP Aggregate Address configuration panel.
   - If the device does not have any BGP aggregate addresses configured, the BGP Aggregate Address configuration panel is displayed, as shown in the following example.
   - If a BGP aggregate address is already configured and you are adding a new one, click on the Add Aggregate Address link to display the BGP Aggregate Address configuration panel, as shown in the following example.
   - If you are modifying an existing BGP aggregate address, click on the Modify button to the right of the row describing the aggregate address to display the BGP Aggregate Address configuration panel, as shown in the following example.

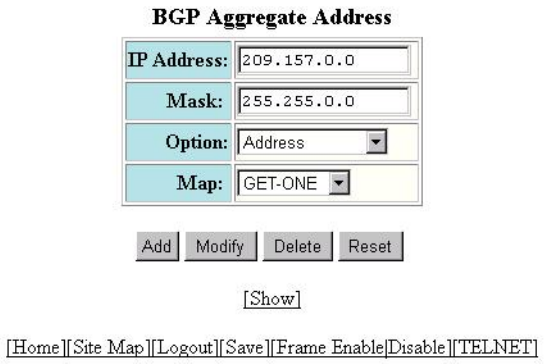

5. Enter the aggregate address in the IP Address field. Specify 0 for the host portion and for the network portion that differs among the networks in the aggregate. For example, to aggregate 10.0.1.0, 10.0.2.0, and 10.0.3.0, enter the IP address 10.0.0.0. Then enter 255.255.0.0 in the Mask field.

6. Enter the mask in the Mask field.
7. 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 router to aggregate AS-path information for all the routes in the aggregate address into a single AS-path.
   - Summary Only – This option prevents the router 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 router to advertise the more specific routes in the specified route map.
   - Attribute Map – This option configures the router to set attributes for the aggregate routes based on the specified route map.
8. 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 8-55 for information on defining a route map.

9. Click the Add button to apply the changes to the device's running-config file.
10. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 Routes by Route Type

You can easily configure BGP4 to redistribute routes of a specific route type using the following methods.

USING THE CLI

To enable redistribution of all OSPF routes and directly attached routes into BGP4, enter the following commands.

```
HP9300(config)# router bgp
HP9300(config-bgp-router)# redistribution ospf
HP9300(config-bgp-router)# redistribution connected
HP9300(config-bgp-router)# write memory
```

Syntax: [no] redistribution connected | ospf | rip | static

USING THE WEB MANAGEMENT INTERFACE

Use the procedure in "Redistributing RIP Routes" on page 8-37.

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 8-55 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, static, and connected (directly attached) routes.

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Redistribute link to display the BGP Redistribute configuration panel.
   - If the device does not have any BGP redistribution parameters configured, the BGP Redistribute configuration panel is displayed, as shown in the following example.
   - If BGP redistribution parameters are already configured and you are adding new ones, click on the Add Redistribute link to display the BGP Redistribute configuration panel, as shown in the following example.
   - If you are modifying existing BGP redistribution parameters, click on the Modify button to the right of the row describing the redistribution parameters to display the BGP Redistribute configuration panel, as shown in the following example.

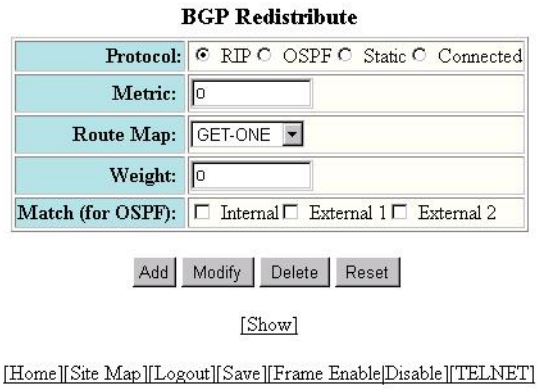

5. Select the source of the routes you want to redistribute into BGP4. You can select RIP, OSPF, Static, or Connected (directly attached).
6. Optionally enter a metric for the redistributed routes in the Metric field. You can specify a value from 0 – 4294967295. The default is 0.
7. Optionally select a route map from the Map field’s pulldown list.

NOTE: The route map must already be defined. See “Defining Route Maps” on page 8-55 for information on defining a route map.

---

8. Optionally enter a weight for the redistributed routes in the Weight field. You can specify a value from 0 – 65535. The default is 0.

9. For OSPF routes, select one of the following to specify the types of OSPF routes to be redistributed into BGP4:
   - Internal
   - External 1
   - External 2
10. Click the Add button to apply the changes to the device's running-config file.
11. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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 8-55 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 8-37.

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 8-55 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 8-37.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Address Filter link to display the BGP Address Filter panel.
    - If the device does not have any BGP address filters configured, the BGP Address Filter configuration panel is displayed, as shown in the following example.
    - If BGP address filters are already configured and you are adding a new one, click on the Add Address Filter link to display the BGP Address Filter configuration panel, as shown in the following example.
    - If you are modifying an existing BGP address filter, click on the Modify button to the right of the row describing the filter to display the BGP Address Filter configuration panel, as shown in the following example.

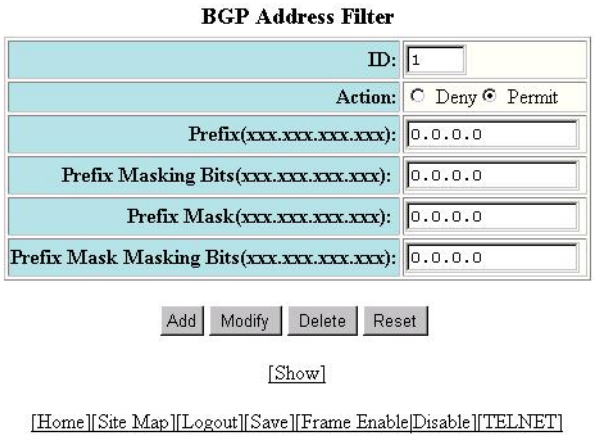

5. Enter the filter ID in the ID field. You can specify a number from 1 – 100.
6. Select the action you want the routing switch to perform if the filter is true:
   - If you select Deny, the router denies the route from entering the BGP4 table if the filter match is true.
   - If you select Permit, the router permits the route into the BGP4 table if the filter match is true.
7. Enter the network prefix in the Prefix field. If you specify “any”, all networks match the filter.
8. 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.
9. Enter the mask in the Prefix Mask field. If you specify “any”, all masks match the filter.
10. Enter the masking bits for the network mask in the Prefix Mask Masking Bits field.
11. Click the Add button to apply the changes to the device’s running-config file.
12. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

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.

The routing switch provides the following methods for filtering on AS-path information:

- AS-path filters
- AS-path ACLs

NOTE: The routing switch cannot actively support AS-path filters and AS-path ACLs at the same time. Use one method or the other but do not mix methods.

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

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

Defining an AS-Path Filter

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. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the AS Path Filter link to display the BGP AS Path Filter panel.
   - If the device does not have any BGP AS-path filters configured, the BGP AS Path Filter configuration panel is displayed, as shown in the following example.
   - If BGP AS-path filters are already configured and you are adding a new one, click on the Add AS Path Filter link to display the BGP AS Path Filter configuration panel, as shown in the following example.
   - If you are modifying an existing BGP AS-path filter, click on the Modify button to the right of the row describing the filter to display the BGP AS Path Filter configuration panel, as shown in the following example.

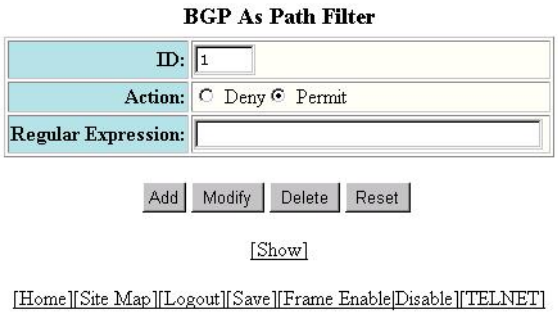

5. Enter the filter ID in the ID field. You can specify a number from 1 – 100.

6. Select the action you want the routing switch to perform if the filter is true:
   - If you select Deny, the router denies the route from entering the BGP4 table if the filter match is true.
   - If you select Permit, the router permits the route into the BGP4 table if the filter match is true.
7. 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 8-44.
8. Click the Add button to apply the changes to the device's running-config file.
9. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Defining an AS-Path ACL

To configure an AS-path ACL, use either of the following methods.

USING THE CLI

To configure an AS-path list that uses ACL 1, enter a command such as the following:

```
HP9300(config)# ip as-path access-list 1 permit 100
HP9300(config)# router bgp
HP9300(config-bgp-router)# neighbor 10.10.10.1 filter-list 1 in
```

The **ip as-path** command configures an AS-path ACL that permits routes containing AS number 100 in their AS paths. The **neighbor** command then applies the AS-path ACL to advertisements and updates received from neighbor 10.10.10.1. In this example, the only routes the routing switch permits from neighbor 10.10.10.1 are those whose AS-paths contain AS-path number 100.

Syntax: ip as-path access-list <num> [seq <seq-value>] deny | permit <as-regular-expression>

The <num> parameter specifies the ACL number and can be from 1 – 199.

The **seq** <seq-value> parameter is optional and specifies the AS-path list's sequence number. You can configure up to 199 entries in an AS-path list. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with number 5. The software interprets the entries in an AS-path list in numerical order, beginning with the lowest sequence number.

The **deny | permit** parameter specifies the action the software takes if a route's AS-path list matches a match statement in this ACL. To configure the AS-path match statements, use the **match as-path** command. See "Matching Based on AS-Path ACL" on page 8-59.

The <as-regular-expression> parameter specifies the AS path information you want to permit or deny to routes that match any of the match statements within the ACL. You can enter a specific AS number or use a regular expression. For the regular expression syntax, see "Using Regular Expressions" on page 8-44.

The **neighbor** command uses the **filter-list** parameter to apply the AS-path ACL to the neighbor. See "Adding BGP4 Neighbors" on page 8-14.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the list of configuration options.
3. Click on the plus sign next to IP to display the list of IP configuration options.
4. Select the AS Path Access List link.
   - If the device does not have any AS Path ACLs, the IP AS Path Access List panel is displayed, as shown in the following example.
   - If an AS Path ACL is already configured and you are adding a new one, click on the Add AS Path Access List link to display the IP AS Path Access List panel, as shown in the following example.

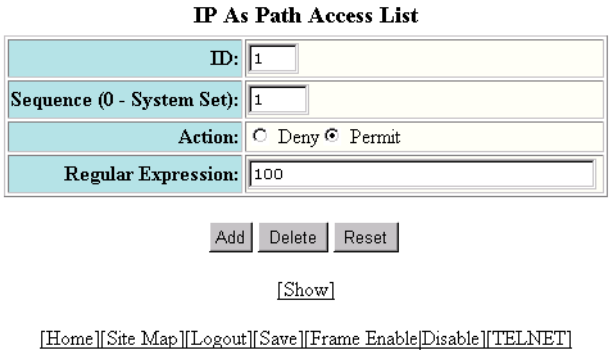

NOTE: You cannot modify an AS Path ACL. Instead, you can delete and then re-add the ACL. To delete an ACL, click on the Delete button to the right of the row describing the ACL, then click on the Add AS Path Access List link.

5. Edit the ACL ID in the ID field, if needed. You can enter a number from 1 – 199.
6. Edit the number in the sequence number in the Sequence field, if you want to override the automatically generated sequence number. You can configure up to 199 entries in an AS-path list. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with number 5. The software interprets the entries in an AS-path list in numerical order, beginning with the lowest sequence number.
7. Select the action you want the software to perform if a route's AS path list matches this ACL entry. You can select Deny or Permit.
8. Enter a regular expression to specify the AS path information you want to permit or deny to routes that match this ACL entry. You can enter a specific AS number or use a regular expression. For the regular expression syntax, see "Using Regular Expressions" on page 8-44.
9. Click the Add button to save the change to the device's running-config file.
10. Repeat steps 6 – 9 for each entry in the ACL. To create another AS Path ACL, go to step 5.
11. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You cannot apply the AS path ACLs to a neighbor using the Web management interface. You must use the CLI. The AS Path Filter List for Weight field in the BGP Neighbor panel of the Web management interface is not used for AS path filtering, but is instead used for changing a route's weight based on the AS path list.

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 8.8.1 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 8.1: BGP4 Special Characters for Regular Expressions

Character	Operation
.	The period matches on any single character, including a blank space. For example, the following regular expression matches for “aa”, “ab”, “ac”, and so on, but not just “a”. a.
*****	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*
+	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 a sequence of “g”s, such as “deg”, “degg”, “deggg”, and so on: 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
^	A caret (when not used within brackets) matches on the beginning of an input string. For example, the following regular expression matches on an AS-path that begins with “jlampa”: ^jlampa
$	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$

Table 8.1: BGP4 Special Characters for Regular Expressions (Continued)

Character	Operation
_	An underscore matches on one or more 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 "100" but not on "1002", "2100", and so on. _100_
[ ]	Square brackets enclose a range of single-character patterns. For example, the following regular expression matches on an AS-path that contains "1", "2", "3", "4", or "5": [1-5] You can use the following expression symbols within the brackets. These symbols are allowed only inside the brackets. • **^ –** The caret matches on any characters *except* the ones in the brackets. For example, the following regular expression matches on an AS-path that does *not* contain "1", "2", "3", "4", or "5": [^1-5] • **-** The hyphen separates the beginning and ending of a range of characters. A match occurs if any of the characters within the range is present. See the example above.
|	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 "abc" or "defg": (abc)|(defg) **Note**: The parentheses group multiple characters to be treated as one value. See the following row for more information about parentheses.
( )	Parentheses allow you to create complex expressions. For example, the following complex expression matches on "abc", "abcabc", or "abcabcabcdefg", but not on "abcdefgdefg": ((abc)+)|((defg)?)

If you want to filter for a special character instead of using the special character as described in Table 8.8.1, 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

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

A community is an optional attribute that identifies the route as a member of a user-defined class of routes. Community names are arbitrary values made of two five-digit integers joined by a colon. You determine what the name means when you create the community name as one of a route's attributes. Each string in the community name can be a number from 0 – 65535.

This format allows you to easily classify community names. For example, a common convention used in community naming is to configure the first string as the local AS and the second string as the unique community within that AS. Using this convention, communities 1:10, 1:20, and 1:30 can be easily identified as member communities of AS 1.

The routing switch provides the following methods for filtering on AS-path information:

- Community filters
- Community list ACLs

NOTE: The routing switch cannot actively support community filters and community list ACLs at the same time. Use one method or the other but do not mix methods.

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

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

Defining a Community Filter

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 <num>:<num> | internet | local-as | 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 router applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the router 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 router takes if the filter match is true.

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

The <num>:<num> parameter indicates a specific community number to filter. Use this parameter to filter for a private (administrator-defined) community. You can enter up to 20 community numbers with the same command.

If you want to filter for the well-known communities “LOCAL_AS”, “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 **local-as** keyword checks for routes with the well-known community “LOCAL_AS”. This community applies only to confederations. The routing switch advertises the route only within the sub-AS. For information about confederations, see “Configuring Confederations” on page 8-33.

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. If the router is a member of a confederation, the routing switch advertises the route only within the confederation. For information about confederations, see “Configuring Confederations” on page 8-33.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Community Filter link to display the BGP Community Filter panel.

NOTE: If the device already has community filters, a table listing the filters is displayed. Click the Modify button to the right of the row describing a filter to change its configuration, or select the Add Community Filter link to display the BGP Community Filter panel.

---

5. Enter the filter’s position in the ID filter. The ID is 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 router applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the router stops and does not continue applying filters from the list.

   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.
6. Select the action for the filter. You can select Deny or Permit:
   - If you select Deny, the router denies the route from entering the BGP4 table if the filter match is true.
   - If you select Permit, the router permits the route into the BGP4 table if the filter match is true.
7. Specify a well-known community you want the routing switch to apply to a route when the route matches the filter by selecting from the following:
   - 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.
   - Local AS – Filters for routes with the well-known community “LOCAL_AS”. A route in this community should not be advertised outside the sub-AS. This community type applies to confederations.
   - No Advertise – 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. If the router is a member of a confederation, the routing switch advertises the route only within the confederation.

NOTE: If you want to filter on a private (administrator-defined) community, do not select one of these. Instead, enter the community number in the Community List field.

---

8. Specify private communities by entering the community names in the Community List field. Enter the names in the following format <num>:<num>. You can use commas or spaces to separate the names.
9. Click the Add button (if you are adding a new filter) or the Modify button (if you are changing a filter) to apply the changes to the device's running-config file.
10. Select the Save link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Defining a Community ACL

To configure a community ACL, use either of the following methods.

USING THE CLI

To configure community ACL 1, enter a command such as the following:

```
HP9300(config)# ip community-list 1 permit 123:2
```

This command configures a community ACL that permits routes that contain community 123:2.

NOTE: See "Matching Based on Community ACL" on page 8-60 for information about how to use a community list as a match condition in a route map.

Syntax: ip community-list <num> [seq <seq-value>] deny | permit <community-num>

The <num> parameter specifies the ACL number and can be from 1 – 199.

The **seq** <seq-value> parameter is optional and specifies the community list's sequence number. You can configure up to 199 entries in a community list. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with number 5. The software interprets the entries in a community list in numerical order, beginning with the lowest sequence number.

The **deny | permit** parameter specifies the action the software takes if a route's community list matches a match statement in this ACL. To configure the community-list match statements, use the **match community-list** command. See "Matching Based on Community ACL" on page 8-60.

The <community-num> parameter specifies the community type or community number. This parameter can have the following values:

- <num>**:**<num> – A specific community number
- **internet** – The Internet community
- **no-export** – The community of sub-ASs within a confederation. Routes with this community can be exported to other sub-ASs within the same confederation but cannot be exported outside the confederation to other ASs or otherwise sent to EBGP neighbors.
- **local-as** – The local sub-AS within the confederation. Routes with this community can be advertised only within the local subAS.
- **no-advertise** – Routes with this community cannot be advertised to any other BGP4 routers at all.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the list of configuration options.
3. Click on the plus sign next to IP to display the list of IP configuration options.
4. Select the Community Access List link.
   - If the device does not have any community ACLs, the IP Community List panel is displayed, as shown in the following example.
   - If a community ACL is already configured and you are adding a new one, click on the Add Community Access List link to display the IP Community List panel, as shown in the following example.

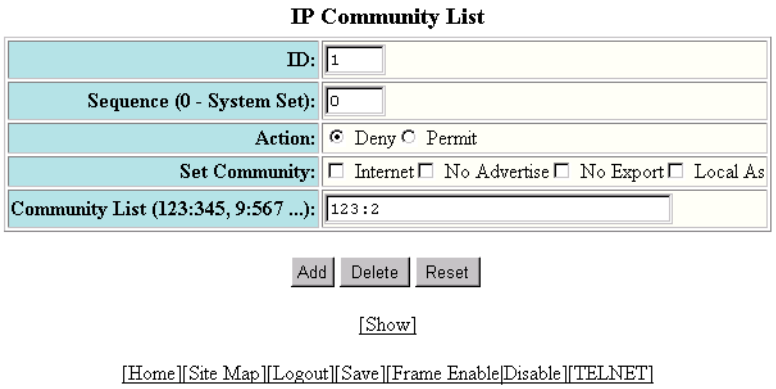

NOTE: You cannot modify a community ACL. Instead, you can delete and then re-add the ACL. To delete an ACL, click on the Delete button to the right of the row describing the ACL, then click on the Add Community List link.

5. Edit the ACL ID in the ID field, if needed. You can enter a number from 1 – 199.
6. Edit the number in the sequence number in the Sequence field, if you want to override the automatically generated sequence number. You can configure up to 199 entries in a community list. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with number 5. The software interprets the entries in ascending sequence order.
7. Select the action you want the software to perform if a route's community list matches this ACL entry.
8. Select the community type by clicking on the checkbox to the left of the description, or enter the community numbers in the Community List field.
9. Click the Add button to save the change to the device's running-config file.
10. Repeat steps 6 – 9 for each entry in the ACL. To create another community ACL, go to step 5.
11. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You cannot apply the community list ACLs to a neighbor using the Web management interface. You must use the CLI.

Defining IP Prefix Lists

An IP prefix list specifies a list of networks. When you apply an IP prefix list to a neighbor, the routing switch sends or receives only a route whose destination is in the IP prefix list. You can configure up to 100 prefix lists. The software interprets the prefix lists in order, beginning with the lowest sequence number.

USING THE CLI

To configure an IP prefix list and apply it to a neighbor, enter commands such as the following:

```
HP9300(config)# ip prefix-list Routesfor20 permit 20.20.0.0/24
HP9300(config-bgp-router)# neighbor 10.10.10.1 prefix-list Routesfor20 out
```

These commands configure an IP prefix list named Routesfor20, which permits routes to network 20.20.0.0. The **neighbor** command configures the routing switch to use IP prefix list Routesfor20 to determine which routes to send to neighbor 10.10.10.1. The routing switch sends routes that go to 20.20.x.x to neighbor 10.10.10.1 because the IP prefix list explicitly permits these routes to be sent to the neighbor.

Syntax: ip prefix-list <name> [seq <seq-value>] [description <string>] deny | permit <network-addr>/<mask-bits> [ge <ge-value>] [le <le-value>]

The <name> parameter specifies the prefix list name. You use this name when applying the prefix list to a neighbor.

The **description** <string> parameter is a text string describing the prefix list.

The **seq** <seq-value> parameter is optional and specifies the IP prefix list's sequence number. You can configure up to 100 prefix list entries. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with prefix list entry 5. The software interprets the prefix list entries in numerical order, beginning with the lowest sequence number.

The **deny | permit** parameter specifies the action the software takes if a neighbor's route is in this prefix list.

The prefix-list matches only on this network unless you use the **ge** <ge-value> or **le** <le-value> parameters. (See below.)

The <network-addr>/<mask-bits> parameter specifies the network number and the number of bits in the network mask.

The **ge** <ge-value> parameter specifies a range of prefixes more specific than the range specified by <network-addr>/<mask-bits> that still match the prefix list. The <ge-value> specifies the minimum number of mask bits in the network mask. For example, if you add **ge 24** to the command above, the prefix list matches on all network numbers that are equal to or more specific than 20.20.0.0. Thus 20.20.1.0 and higher also match the prefix list.

The **le** <ge-value> parameter specifies a range of prefixes less specific than the range specified by <network-addr>/<mask-bits> that still match the prefix list. The <le-value> specifies the maximum number of bits in the mask. For example, if you add **le 16** to the command above, the prefix list matches on 20.20.x.x and on all other 20.x.x.x networks.

If you do not specify **ge** <ge-value> or **le** <le-value>, the prefix list matches only on the exact network prefix you specify with the <network-addr>/<mask-bits> parameter.

For the syntax of the **neighbor** command shown in the example above, see "Adding BGP4 Neighbors" on page 8-14.

USING THE WEB MANAGEMENT INTERFACE

To configure an IP Prefix List, use the following procedure:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the list of configuration options.
3. Click on the plus sign next to IP to display the list of IP configuration options.
4. Select the Prefix List link.
    - If the device does not have any prefix list ACLs, the IP Prefix List panel is displayed, as shown in the following example.
    - If a prefix list ACL is already configured and you are adding a new one, click on the Add IP Prefix List link to display the IP Prefix List panel, as shown in the following example.

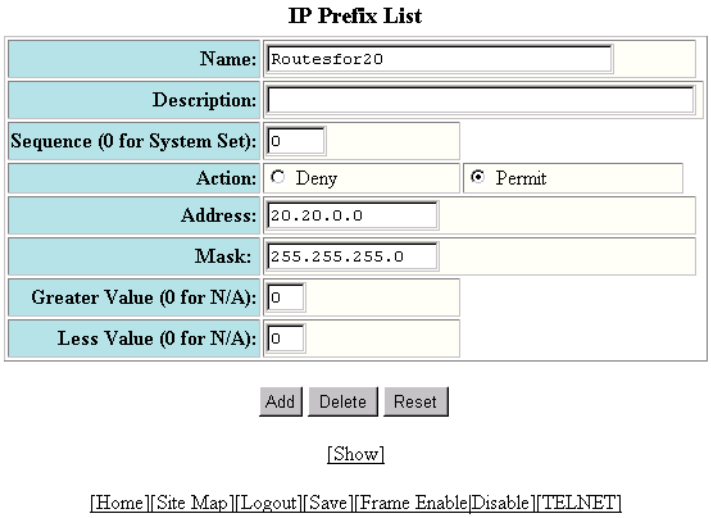

NOTE: You cannot modify an IP prefix list ACL. Instead, you can delete and then re-add the ACL. To delete an ACL, click on the Delete button to the right of the row describing the ACL, then click on the Add IP Prefix List link.

---

5. Edit a name in the Name field.
6. Enter a description in the Description field.
7. Edit the number in the sequence number in the Sequence field, if you want to override the automatically generated sequence number. You can configure up to 100 prefix list entries. If you do not specify a sequence number, the software numbers them in increments of 5, beginning with prefix list entry 5. The software interprets the prefix list entries in numerical order, beginning with the lowest sequence number.
8. Select the action you want the software to perform if a neighbor's route is in this prefix list.
9. Enter the IP prefix by entering a network address and sub-net mask in the Address and Mask fields.

NOTE: If you do not specify a Greater Value or Less Value, this prefix list entry matches only on the exact network prefix you specified with the values in the Address and Mask fields.

---

10. Enter a number from 1 – 32 in the Greater Value field if you want the prefix list to match on prefixes that are more specific than the one you entered in the Address and Mask fields, in addition to matching on the prefix in those fields. The value you enter here specifies the minimum number of mask bits in the network mask. For example, if you enter 24 in the example panel shown above, the prefix list matches on all network numbers that are equal to or more specific than 20.20.0.0. Thus 20.20.1.0 and higher also match the prefix list.
11. Enter a number from 1 – 32 in the Less Value field if you want the prefix list to match on prefixes that are less specific than the one you entered in the Address and Mask fields, in addition to matching on the prefix in those fields.
12. Click the Add button to save the change to the device's running-config file.
13. Repeat steps 5 – 12 for each IP prefix list entry.
14. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

To apply the IP Prefix List to a neighbor, use the following procedure:

1. In the tree view, click on the plus sign next to BGP under Configure to display the list of BGP configuration options.
2. Select the Neighbor link to display the BGP Neighbor panel.
3. Select the Prefix List link to display the BGP Neighbor Prefix List panel, as shown in the following example.

4. Select the neighbor's IP address from the IP Address field's pulldown menu.

NOTE: The address appears in this menu only if you have already configured the neighbor information on the routing switch.

---

5. Select the direction to which you are applying the prefix list by clicking next to In or Out.
   - In – The prefix list applies to routes received from the neighbor.
   - Out – The prefix list applies to routes destined to be sent to the neighbor.
6. Enter the prefix list name or ID in the Prefix List Name field.
7. Click the Add button to save the change to the device's running-config file.
8. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Defining Neighbor Distribute Lists

A neighbor distribute list is a list of BGP4 address filters or ACLs that filter the traffic to or from a neighbor. To configure a neighbor distribute list, use either of the following methods.

USING THE CLI

To configure a distribute list that uses ACL 1, enter a command such as the following:

```
HP9300(config-bgp-router)# neighbor 10.10.10.1 distribute-list 1 in
```

This command configures the routing switch to use ACL 1 to select the routes that the routing switch will accept from neighbor 10.10.10.1.

Syntax: neighbor <ip-addr> distribute-list <name-or-num> in | out

The <ip-addr> parameter specifies the neighbor.

The <name-or-num> parameter specifies the name or number of a standard, extended, or named ACL.

The **in | out** parameter specifies whether the distribute list applies to inbound or outbound routes:

- **in** – controls the routes the routing switch will accept from the neighbor.
- **out** – controls the routes sent to the neighbor.

NOTE: The command syntax shown above is new in software release 06.x. However, the **neighbor** <ip-addr> **distribute-list in | out** <num> command (where the direction is specified before the filter number) is the same as in earlier software releases. Use the new syntax when you are using an IP ACL or IP prefix list with the distribute list. Use the old syntax when you are using a BGP4 address filter with the distribute list.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Neighbor link to display the BGP Neighbor panel.

NOTE: If the device already has neighbors, a table listing the neighbors is displayed. Click the Modify button to the right of the row describing the neighbor to change its configuration, or click the Add Neighbor link to display the BGP Neighbor configuration panel.

5. If you are adding a new neighbor or you need to change additional parameters, see the complete procedure in “Adding BGP4 Neighbors” on page 8-14.
6. Select the Distribute List link at the bottom of the panel to display the BGP Neighbor Distribute panel, as shown in the following example.

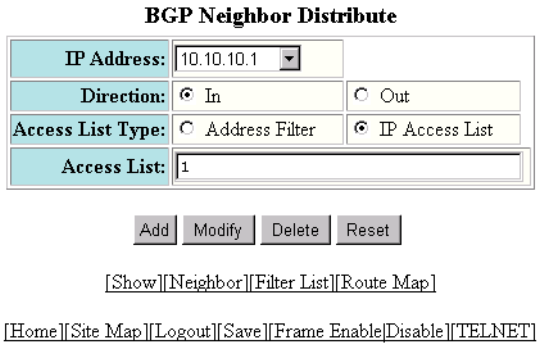

7. Select the neighbor’s IP address from the IP Address field’s pulldown menu.

NOTE: The address appears in this menu only if you have already configured the neighbor information on the routing switch.

8. Select the direction to which you are applying the distribute list by clicking next to In or Out.
   - In – The distribute list applies to routes received from the neighbor.
   - Out – The distribute list applies to routes destined to be sent to the neighbor.
9. Select the type of distribute list you are applying. You can select one of the following:
   - Address Filter – a BGP4 address filter.
   - IP Access List – an ACL.
10. Enter the address filter or ACL name or ID in the Access List field.
11. Click the Add button to save the change to the device’s running-config file.
12. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Defining Route Maps

A ***route map*** is a named set of match conditions and parameter settings that the router 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 router 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 router 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.
- 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 router 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 router 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 8-57 and "Setting Parameters in the Routes" on page 8-61.

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 router.

The **permit | deny** parameter specifies the action the router 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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Route Map Filter link to display the BGP Route Map Filter panel.
    - If the device does not have any BGP route map filters configured, the BGP Route Map Filter configuration panel is displayed, as shown in the following example.
    - If BGP route map filters are already configured and you are adding a new one, click on the Route Map Filter link to display the BGP Route Map Filter configuration panel, as shown in the following example.
    - If you are modifying an existing BGP route map filter, click on the Modify button to the right of the row describing the filter to display the BGP Route Map Filter configuration panel, as shown in the following example.

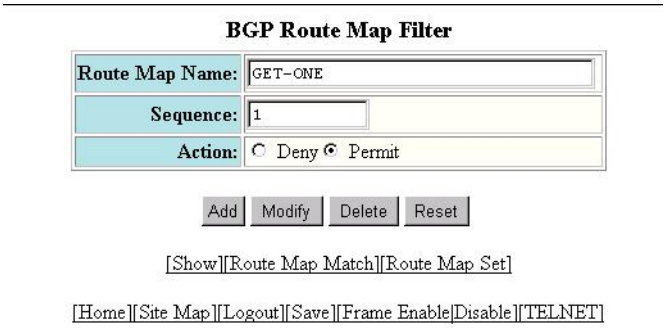

5. Enter the name of the route map in the Route Map Name field.
6. Enter the sequence (instance) number in the Sequence field. The routing switch applies the instances in ascending numerical order. Once an instance comparison results in a “true” evaluation, the routing switch stops applying instances and applies the match and set statements you configure for the instance. See “Specifying the Match Conditions” on page 8-57 and “Setting Parameters in the Routes” on page 8-61.
7. 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.
8. Click the Add button to apply the changes to the device’s running-config file.
9. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

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 <name-or-num>] |
[address-filters | as-path-filters | community-filters | community-list <num,num,...>] |
[ip address | next-hop <acl-name-or-num> | prefix-list <string>] | [metric <num>] |
[next-hop <address-filter-list>] | [route-type internal | external-type1 | external-type2] | [tag <tag-value>]

The **as-path** <name-or-num> parameter specifies an AS-path ACL. You can specify up to five AS-path ACLs. To configure an AS-path ACL, use the **ip as-path access-list** command. See “Defining an AS-Path ACL” on page 8-43.

The **address-filters | as-path-filters | community-filters | community-list** <num,num,...> parameter specifies a filter or list of filters to be matched for each route. The router treats the first match as the best match. If a route does not match any filter in the list, then the router considers the match condition to have failed. To configure an address filter, see “Filtering Specific IP Addresses” on page 8-40. To configure an AS-path filter or AS-path ACL, see “Filtering AS-Paths” on page 8-41. To configure a community filter or community ACL, see “Filtering Communities” on page 8-47.

You can enter up to six community names on the same command line.

NOTE: The filters or ACLs must already be configured.

The **ip address | next-hop** <acl-name-or-num> **| prefix-list** <string> parameter specifies an ACL or IP prefix list. Use this parameter to match based on the destination network or next-hop gateway. To configure an IP ACL for use with this command, use the **ip access-list** command. See “Using Access Control Lists (ACLs)” on page 3-1. To configure an IP prefix list, use the **ip prefix-list** command. See “Defining IP Prefix Lists” on page 8-50.

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. If you have just added the route map and the map is displayed in the BGP Route Map Filter panel, go to step 7. Otherwise, go to step 2.
2. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
3. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
4. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
5. Click on the Route Map Filter link to display a table listing the configured BGP route maps.
6. Click Modify next to the route map you want to configure to display the map in the BGP Route Map Filter panel.
7. Select the Route Map Match link at the bottom of the panel to display the BGP Route Map Match panel.
8. 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.
9. For OSPF routes, select the one of the following route types—Internal, External1, or External2.
10. Select the type of ACL or filter you are adding as a match condition. You can select more than one ACL or filter type. In this example, select AS Path Access List.

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

    NOTE: The routing switch does not actively support both filters and ACLs at the same time. Use one method or the other.

    NOTE: IP prefix lists and neighbor distribute lists provide separate means for the same type of filtering. To simplify configuration, HP recommends you use one method or the other but do not mix them.

    ---

11. Enter the filter or ACL numbers or names in the entry fields next to the filter or ACL types you selected.
12. 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.
13. 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.
14. 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.
15. Click the Apply button to apply the changes to the device's running-config file.
16. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Match Examples Using ACLs

The following sections show some detailed examples of how to configure route maps that include match statements that match on ACLs.

Matching Based on AS-Path ACL

To construct match statements for a route map that match based on AS-path information, use either of the following methods.

USING THE CLI

To construct a route map that matches based on AS-path ACL 1, enter the following commands:

```
HP9300(config)# route-map PathMap permit 1
HP9300(config-bgp-routemap PathMap)# match as-path 1
```

Syntax: match as-path <name-or-num>

The <name-or-num> parameter specifies an AS-path ACL and can be a number from 1 – 199. You can specify up to five AS-path ACLs. To configure an AS-path ACL, use the **ip as-path access-list** command. See “Defining an AS-Path ACL” on page 8-43.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Route Map Filter link to display the BGP Route Map Filter panel.
   - If the device does not have any BGP route map filters configured, the BGP Route Map Filter configuration panel is displayed, as shown in the following example.
   - If BGP route map filters are already configured and you are adding a new one, click on the Route Map Filter link to display the BGP Route Map Filter configuration panel, as shown in the following example.
   - If you are modifying an existing BGP route map filter, click on the Modify button to the right of the row describing the filter to display the BGP Route Map Filter configuration panel, as shown in the following example.

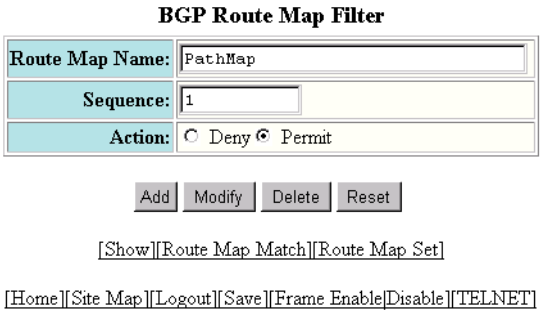

5. Enter the name of the route map in the Route Map Name field.
6. Enter the sequence (instance) number in the Sequence field. The routing switch applies the instances in ascending numerical order. Once an instance comparison results in a “true” evaluation, the routing switch stops applying instances and applies the match and set statements you configure for the instance.
7. 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.

8. Click the Add button to apply the changes to the device's running-config file.
9. Select the Route Map Match link at the bottom of the panel to display the BGP Route Map Match panel, as shown in the following example.

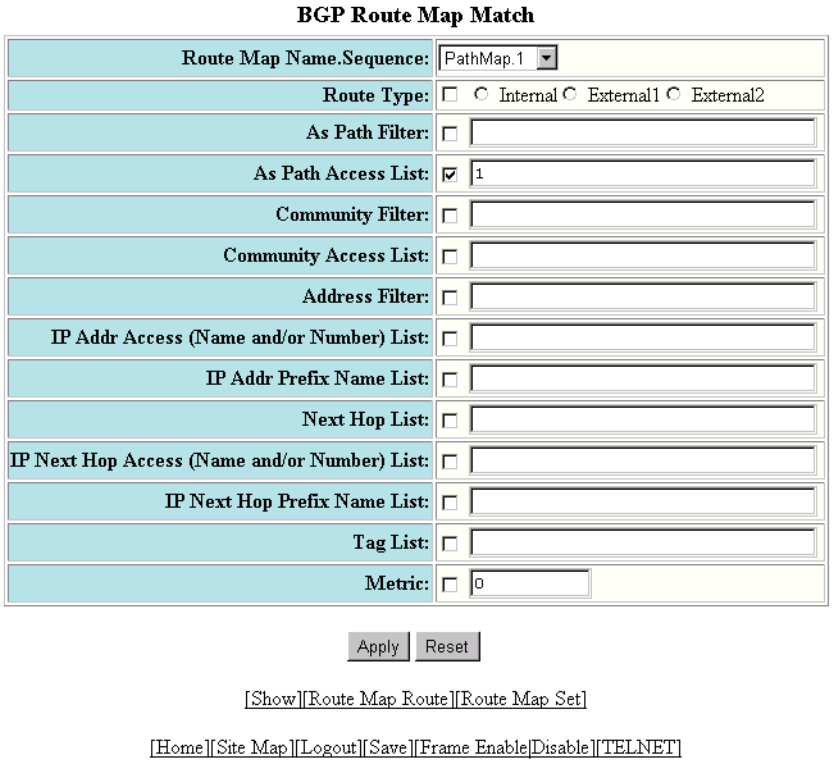

10. Select the type of ACL or filter you are adding as a match condition. You can select more than one ACL or filter type. In this example, select AS Path Access List.

NOTE: IP prefix lists and neighbor distribute lists provide separate means for the same type of filtering. To simplify configuration, HP recommends you use one method or the other but do not mix them.

---

11. Next to each type of ACL or filter you selected, enter the ACL or filter name or ID. In this example, AS-path ACL 1 is specified.
12. Click the Apply button to save the change to the device's running-config file.
13. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Matching Based on Community ACL

To construct match statements for a route map that match based on community information, use either of the following methods.

USING THE CLI

To construct a route map that matches based on community ACL 1, enter the following commands:

```
HP9300(config)# ip community-list 1 permit 123:2
HP9300(config)# route-map CommMap permit 1
HP9300(config-bgp-routemap CommMap)# match community-list 1
```

Syntax: match community-list <name-or-num>

The <name-or-num> parameter specifies a community list ACL and can be a number from 1 – 199. To configure a community list ACL, use the **ip community-list** command. See "Defining a Community ACL" on page 8-49.

USING THE WEB MANAGEMENT INTERFACE

Use the procedure in “Matching Based on AS-Path ACL” on page 8-59, but select Community Access List instead of AS Path Access List.

Matching Based on Destination Network

To construct match statements for a route map that match based on destination network, use either of the following methods. You can use the results of an IP ACL or an IP prefix list as the match condition.

USING THE CLI

To construct a route map that matches based on destination network, enter commands such as the following:

```
HP9300(config)# route-map NetMap permit 1
HP9300(config-bgp-routemap NetMap)# match ip address 1
```

Syntax: match ip address <name-or-num>

Syntax: match ip address prefix-list <name>

The <name-or-num> parameter with the first command specifies an IP ACL and can be a number from 1 – 199 or the ACL name if it is a named ACL. To configure an IP ACL, use the **ip access-list** or **access-list** command. See “Using Access Control Lists (ACLs)” on page 3-1.

The <name> parameter with the second command specifies an IP prefix list name. To configure an IP prefix list, see “Defining IP Prefix Lists” on page 8-50.

USING THE WEB MANAGEMENT INTERFACE

Use the procedure in “Matching Based on AS-Path ACL” on page 8-59, but select IP Addr Access (Name and/or Number) List instead of AS Path Access List.

Matching Based on Next-Hop Router

To construct match statements for a route map that match based on the IP address of the next-hop router, use either of the following methods. You can use the results of an IP ACL or an IP prefix list as the match condition.

USING THE CLI

To construct a route map that matches based on the next-hop router, enter commands such as the following:

```
HP9300(config)# route-map HopMap permit 1
HP9300(config-bgp-routemap HopMap)# match ip next-hop 2
```

Syntax: match ip next-hop <name-or-num>

Syntax: match ip next-hop prefix-list <name>

The <name-or-num> parameter with the first command specifies an IP ACL and can be a number from 1 – 199 or the ACL name if it is a named ACL. To configure an IP ACL, use the **ip access-list** or **access-list** command. See “Using Access Control Lists (ACLs)” on page 3-1.

The <name> parameter with the second command specifies an IP prefix list name. To configure an IP prefix list, see “Defining IP Prefix Lists” on page 8-50.

USING THE WEB MANAGEMENT INTERFACE

Use the procedure in “Matching Based on AS-Path ACL” on page 8-59, but select IP Next Hop Access (Name and/or Number) List instead of AS Path Access List.

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 local preference 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 8-63.

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. If you have just added the route map and the map is displayed in the BGP Route Map Filter panel, go to step 7. Otherwise, go to step 2.
2. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
3. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
4. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
5. Click on the Route Map Filter link to display a table listing the configured BGP route maps.
6. Click Modify next to the route map you want to configure to display the map in the BGP Route Map Filter panel.
7. Select the Route Map Set link at the bottom of the panel to display the BGP Route Map Set panel.
8. Select the sequence (instance) from the Route Map Name Sequence field's pulldown list.
9. Optionally select the origin. You can select IGP or Incomplete. Also select the checkbox in front of the field.
10. Optionally enter AS numbers to append to the AS path. Also select the checkbox in front of the field.
11. Optionally select Auto Tag. The routing switch calculates and sets an automatic tag value for the route.
12. 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.
13. Optionally select the community type and also select the checkbox.

14. For a private community, enter the community number in the Number field. You can enter more than one community. Use commas or spaces to separate the community names.
15. Select Additive of you want the Set statement to add the specified community.
16. 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.
17. 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.
18. Optionally enter the Next Hop IP address in the NextHop field and also select the checkbox in front of the field.
19. 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.
20. Click the Apply button to apply the changes to the device's running-config file.
21. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Use the Web management procedures in "Defining Route Maps" on page 8-55 to create the route map.
3. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
4. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.

5. Click on the General link to display the BGP configuration panel.
6. Select the route map name from the Table Map field's pulldown menu.
7. Click the Apply button to apply the changes to the device's running-config file.
8. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring Route Flap Dampening

A "route flap" is the change in a route's state, from up to down or down to up. When a route's state changes, the state change causes changes in the route tables of the routers that support the route. Frequent changes in a route's state can cause Internet instability and add processing overhead to the routers that support the route.

Route flap dampening is a mechanism that reduces the impact of route flap by changing a BGP4 router's response to route state changes. When route flap dampening is configured, the routing switch suppresses unstable routes until the route's state changes reduce enough to meet an acceptable degree of stability. The HP implementation of route flap dampening is based on RFC 2439.

Route flap dampening is disabled by default. You can enable the feature globally or on an individual route basis using route maps.

NOTE: The routing switch applies route flap dampening only to routes learned from EBGP neighbors.

The route flap dampening mechanism is based on penalties. When a route exceeds a configured penalty value, the routing switch stops using that route and also stops advertising it to other routers. The mechanism also allows a route's penalties to reduce over time if the route's stability improves. The route flap dampening mechanism uses the following parameters:

- Suppression threshold – Specifies the penalty value at which the routing switch stops using the route. Each time a route becomes unreachable or is withdrawn by a BGP4 UPDATE from a neighbor, the route receives a penalty of 1000. By default, when a route has a penalty value greater than 2000, the routing switch stops using the route. Thus, by default, if a route goes down more than twice, the routing switch stops using the route. You can set the suppression threshold to a value from 1 – 20000. The default is 2000.
- Half-life – Once a route has been assigned a penalty, the penalty decreases exponentially and decreases by half after the half-life period. The default half-life period is 15 minutes. The software reduces route penalties every five seconds. For example, if a route has a penalty of 2000 and does not receive any more penalties (it does not go down again) during the half-life, the penalty is reduced to 1000 after the half-life expires. You can configure the half-life to be from 1 - 45 minutes. The default is 15 minutes.
- Reuse threshold – Specifies the minimum penalty a route can have and still be suppressed by the routing switch. If the route's penalty falls below this value, the routing switch un-suppresses the route and can use it again. The software evaluates the dampened routes every ten seconds and un-suppresses the routes that have penalties below the reuse threshold. You can set the reuse threshold to a value from 1 - 20000. The default is 750.
- Maximum suppression time – Specifies the maximum number of minutes a route can be suppressed regardless of how unstable the route has been before this time. You can set the parameter to a value from 1 – 20000 minutes. The default is four times the half-life. When the half-life value is set to its default (15 minutes), the maximum suppression time defaults to 60 minutes.

You can configure route flap dampening globally or for individual routes using route maps. If you configure route flap dampening parameters globally and also use route maps, the settings in the route maps override the global values.

Globally Configuring Route Flap Dampening

To configure route flap dampening globally, use either of the following methods.

USING THE CLI

To enable route flap dampening using the default values, enter the following command:

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

Syntax: dampening [<half-life> <reuse> <suppress> <max-suppress-time>]

The <half-life> parameter specifies the number of minutes after which the route's penalty becomes half its value. The route penalty allows routes that have remained stable for a while despite earlier instability to eventually become eligible for use again. The decay rate of the penalty is proportional to the value of the penalty. After the half-life expires, the penalty decays to half its value. Thus, a dampened route that is no longer unstable can eventually become eligible for use again. You can configure the half-life to be from 1 - 45 minutes. The default is 15 minutes.

The <reuse> parameter specifies how low a route's penalty must become before the route becomes eligible for use again after being suppressed. You can set the reuse threshold to a value from 1 – 20000. The default is 750 (0.75, or three-fourths, of the penalty assessed for a one "flap").

The <suppress> parameter specifies how high a route's penalty can become before the routing switch suppresses the route. You can set the suppression threshold to a value from 1 – 20000. The default is 2000 (two "flaps").

The <max-suppress-time> parameter specifies the maximum number of minutes that a route can be suppressed regardless of how unstable it is. You can set the maximum suppression time to a value from 1 – 20000 minutes. The default is four times the half-life setting. Thus, if you use the default half-life of 15 minutes, the maximum suppression time is 60 minutes.

The following example shows how to change the dampening parameters.

```
HP9300(config-bgp-router)# dampening 10 200 2500 40
```

This command changes the half-life to 20 minutes, the reuse threshold to 200, the suppression threshold to 2500, and the maximum number of minutes a route can be dampened to 40.

NOTE: To change any of the parameters, you must specify all the parameters with the command. If you want to leave some parameters unchanged, enter their default values.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the General link to display the BGP configuration panel.
5. Select (Next 4) Parameters next to Dampening, to indicate that you want to enable dampening. This selection also ensures that when you click Apply, the interface applies changes you make to the dampening parameters in the following four fields.
6. Edit the value in the Dampening Half Life field if you want to change the half life. The half like specifies the number of minutes after which the route's penalty becomes half its value. The route penalty allows routes that have remained stable for a while despite earlier instability to eventually become eligible for use again. The decay rate of the penalty is proportional to the value of the penalty. After the half-life. expires, the penalty decays to half its value. Thus, a dampened route that is no longer unstable can eventually become eligible for use again. You can configure the half-life to be from 1 - 45 minutes. The default is 15 minutes.
7. Edit the value in the Dampening Reuse field if you want to change the dampening reuse parameter. The dampening reuse parameter specifies how low a route's penalty must become before the route becomes

eligible for use again after being suppressed. You can set the reuse threshold to a value from 1 – 20000. The default is 750 (0.75, or three-fourths, of the penalty assessed for a one “flap”).

8. Edit the value in the Dampening Suppress field if you want to change the dampening suppress parameter. The dampening suppress parameter specifies how high a route’s penalty can become before the routing switch suppresses the route. You can set the suppression threshold to a value from 1 – 20000. The default is 2000 (two “flaps”).
9. Edit the value in the Dampening Max Suppress Time field if you want to change the maximum suppression parameter. The maximum suppression parameter specifies the maximum number of minutes that a route can be suppressed regardless of how unstable it is. You can set the maximum suppression time to a value from 1 – 20000 minutes. The default is four times the half-life setting. Thus, if you use the default half-life of 15 minutes, the maximum suppression time is 60 minutes.
10. Click the Apply button to apply the changes to the device’s running-config file.
11. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Using a Route Map To Configure Route Flap Dampening for Specific Routes

Route maps enable you to fine tune route flap dampening parameters for individual routes. To configure route flap dampening parameters using route maps, configure BGP4 address filters for each route you want to set the dampening parameters for, then configure route map entries that set the dampening parameters for those routes. The following sections show examples.

To configure route flap dampening for specific routes, use one of the following methods.

USING THE CLI

To configure address filters and a route map for dampening specific routes, enter commands such as the following:

```
HP9300(config)# router bgp
HP9300(config-bgp-router)# address-filter 9 permit 209.157.22.0 255.255.255.0
255.255.255.0 255.255.255.0
HP9300(config-bgp-router)# address-filter 10 permit 209.157.23.0 255.255.255.0
255.255.255.0 255.255.255.0
HP9300(config-bgp-router)# exit
HP9300(config)# route-map DAMPENING_MAP permit 9
HP9300(config-bgp-routemap DAMPENING_MAP)# match address-filters 9
HP9300(config-bgp-routemap DAMPENING_MAP)# set dampening 10 200 2500 40
HP9300(config-bgp-routemap DAMPENING_MAP)# exit
HP9300(config)# route-map DAMPENING_MAP permit 10
HP9300(config-bgp-routemap DAMPENING_MAP)# match address-filters 10
HP9300(config-bgp-routemap DAMPENING_MAP)# set dampening 20 200 2500 60
HP9300(config-bgp-routemap DAMPENING_MAP)# router bgp
HP9300(config-bgp-router)# dampening route-map DAMPENING_MAP
```

The **address-filter** commands in this example configure two BGP4 address filters, for networks 209.157.22.0 and 209.157.23.0. The first route-map command creates an entry in a route map called “DAMPENING_MAP”. Within this entry of the route map, the **match** command matches based on address filter 9, and the **set** command sets the dampening parameters for the route that matches. Thus, for BGP4 routes to 209.157.22.0, the routing switch uses the route map to set the dampening parameters. These parameters override the globally configured dampening parameters.

The commands for the second entry in the route map (instance 10 in this example) perform the same functions for route 209.157.23.0. Notice that the dampening parameters are different for each route.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Address Filter link to display the BGP Address Filter panel.
   - If the device does not have any BGP address filters configured, the BGP Address Filter configuration panel is displayed, as shown in the following example.
   - If BGP address filters are already configured and you are adding a new one, click on the Add Address Filter link to display the BGP Address Filter configuration panel, as shown in the following example.
   - If you are modifying an existing BGP address filter, click on the Modify button to the right of the row describing the filter to display the BGP Address Filter configuration panel, as shown in the following example.

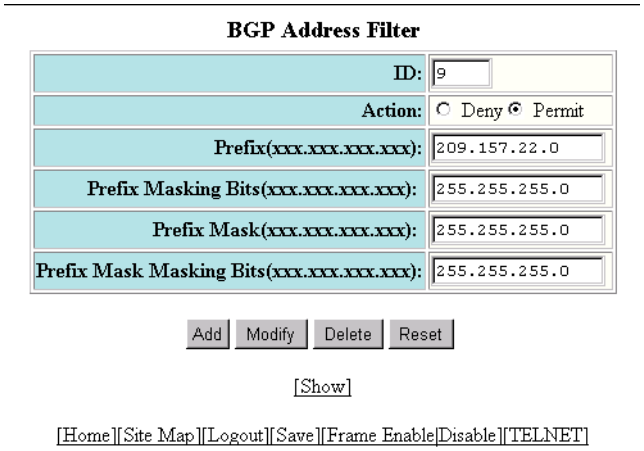

5. Enter the filter ID in the ID field. You can specify a number from 1 – 100.
6. Select the action you want the routing switch to perform if the filter is true:
   - If you select Deny, the router denies the route from entering the BGP4 table if the filter match is true.
   - If you select Permit, the router permits the route into the BGP4 table if the filter match is true.
7. Enter the network prefix in the Prefix field. If you specify “any”, all networks match the filter.
8. 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.
9. Enter the mask in the Prefix Mask field. If you specify “any”, all masks match the filter.
10. Enter the masking bits for the network mask in the Prefix Mask Masking Bits field.
11. Click the Add button to apply the changes to the device’s running-config file.
12. Repeat steps 5 – 11 for each address filter.
13. In the tree view, under BGP in the Configure section, click on the Route Map Filter link to display the BGP Route Map Filter panel.
    - If the device does not have any BGP route map filters configured, the BGP Route Map Filter configuration panel is displayed, as shown in the following example.
    - If BGP route map filters are already configured and you are adding a new one, click on the Route Map Filter link to display the BGP Route Map Filter configuration panel, as shown in the following example.

- If you are modifying an existing BGP route map filter, click on the Modify button to the right of the row describing the filter to display the BGP Route Map Filter configuration panel, as shown in the following example.

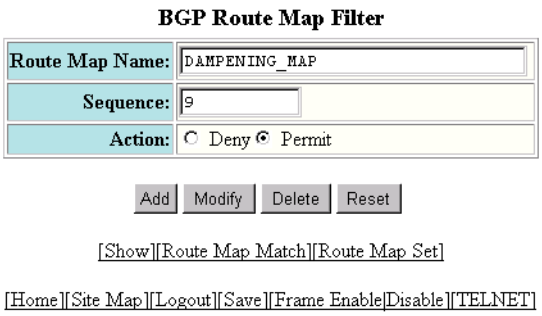

14. Enter the name of the route map in the Route Map Name field.
15. Enter the sequence (instance) number in the Sequence field. The routing switch applies the instances in ascending numerical order. Once an instance comparison results in a “true” evaluation, the routing switch stops applying instances and applies the match and set statements you configure for the instance.

NOTE: In this example, the sequence number matches the address filter number. Using the same number is a convenient way to remember that these configuration items are associated, but is not a requirement.

---

16. 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.
17. Click the Add button to apply the changes to the device’s running-config file.

18. Select the Route Map Match link at the bottom of the panel to display the BGP Route Map Match panel, as shown in the following example.

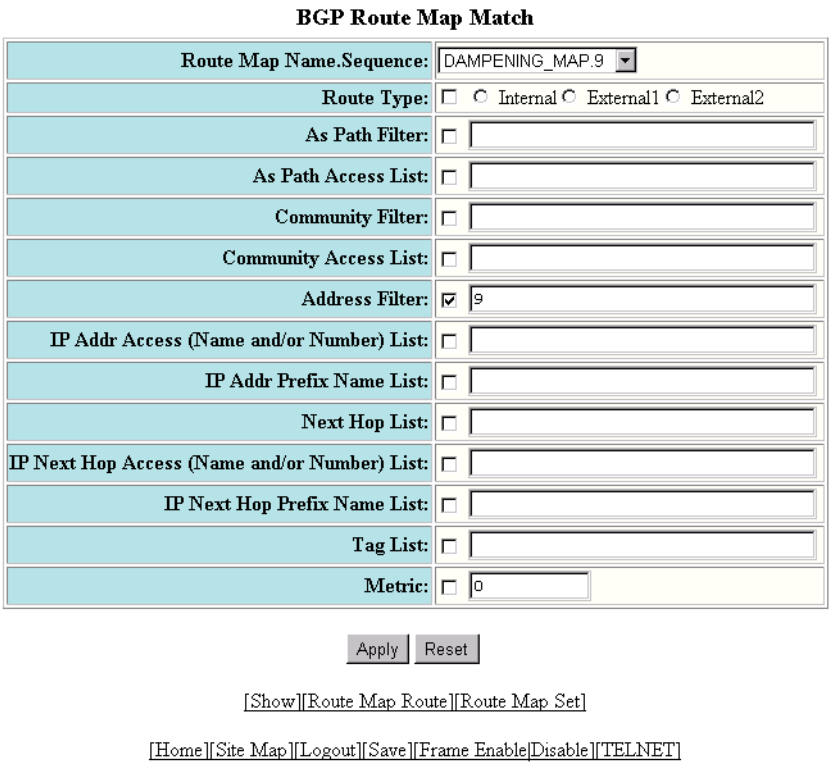

19. Click on the checkbox next to Address Filter to indicate that you are using an address filter as a match condition.
20. Enter the address filter number in the Address Filter field.
21. Click Apply to apply the changes to the device's running-config file.

22. Select the Route Map Set link at the bottom of the panel to display the BGP Route Map Set panel, as shown in the following example.

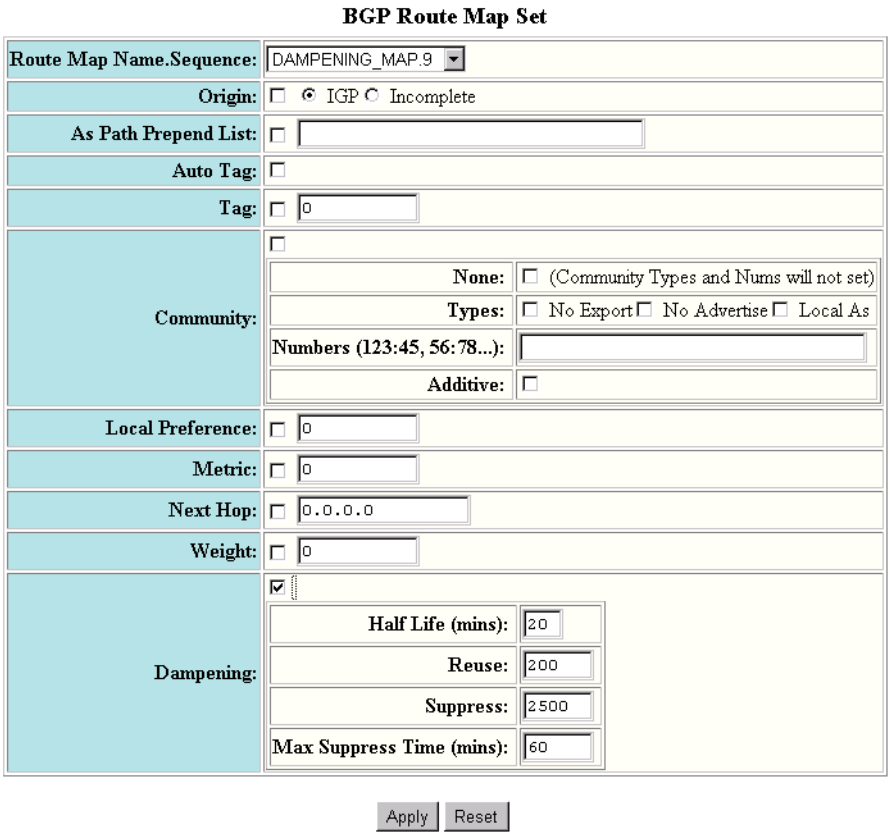

23. Select the checkbox in the Dampening section to specify that this route map is setting dampening parameters.
24. Edit the value in the Half Life field to specify the half life you want this route map to set for routes that match the match conditions you specified above.
25. Edit the value in the Reuse field to specify the dampening reuse value you want this route map to set.
26. Edit the value in the Suppress field to specify the dampening suppress value you want this route map to set.
27. Edit the value in the Max Suppress Time field to specify the maximum suppression value you want this route map to set.
28. Click Apply to apply the changes to the device's running-config file.
29. In the tree view, under BGP in the Configure section, click on the General link to display the BGP configuration panel.
30. In the Dampening section, click next to Route-Map, then select the dampening route map from the Route-Map field's pulldown menu. In this example, select the map named DAMPENING_MAP.

NOTE: The route map appears in this menu only if you have already configured the route map.

31. Click Apply to apply the changes to the device's running-config file.
32. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Using a Route Map To Configure Route Flap Dampening for a Specific Neighbor

You can use a route map to configure route flap dampening for a specific neighbor by performing the following tasks:

- Configure an empty route map with no match or set statements. This route map does not specify particular routes for dampening but does allow you to enable dampening globally when you refer to this route map from within the BGP configuration level.
- Configure another route map that explicitly enables dampening. Use a set statement within the route map to enable dampening. When you associate this route map with a specific neighbor, the route map enables dampening for all routes associated with the neighbor. You also can use match statements within the route map to selectively perform dampening on some routes from the neighbor.

  NOTE: You still need to configure the first route map to enable dampening globally. The second route map does not enable dampening by itself; it just applies dampening to a neighbor.

- Apply the route map to the neighbor.

USING THE CLI

To enable route flap dampening for a specific BGP4 neighbor, enter commands such as the following:

```
HP9300(config)# route-map DAMPENING_MAP_ENABLE permit 1
HP9300(config-bgp-routemap DAMPENING_MAP_ENABLE)# exit
HP9300(config)# route-map DAMPENING_MAP_NEIGHBOR_A permit 1
HP9300(config-bgp-routemap DAMPENING_MAP_NEIGHBOR_A)# set dampening
HP9300(config-bgp-routemap DAMPENING_MAP_NEIGHBOR_A)# exit
HP9300(config)# router bgp
HP9300(config-bgp-router)# dampening route-map DAMPENING_MAP_ENABLE
HP9300(config-bgp-router)# neighbor 10.10.10.1 route-map in
DAMPENING_MAP_NEIGHBOR_A
```

In this example, the first command globally enables route flap dampening. This route map does not contain any match or set statements. At the BGP configuration level, the **dampening route-map** command refers to the DAMPENING_MAP_ENABLE route map created by the first command, thus enabling dampening globally.

The third and fourth commands configure a second route map that explicitly enables dampening. Notice that the route map does not contain a match statement. The route map implicitly applies to all routes. Since the route map will be applied to a neighbor at the BGP configuration level, the route map will apply to all routes associated with the neighbor.

Although the second route map enables dampening, the first route map is still required. The second route map enables dampening for the neighbors to which the route map is applied. However, unless dampening is already enabled globally by the first route map, the second route map has no effect.

The last two commands apply the route maps. The **dampening route-map** command applies the first route map, which enables dampening globally. The **neighbor** command applies the second route map to neighbor 10.10.10.1. Since the second route map does not contain match statements for specific routes, the route map enables dampening for all routes received from the neighbor.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. In the tree view, under BGP in the Configure section, click on the Route Map Filter link to display the BGP Route Map Filter panel.

NOTE: If the device already has route maps, a table listing the route maps is displayed. Click the Modify button to the right of the row describing the route map to change its configuration, or click the Add Route Map Filter link to display the BGP Route Map Filter panel.

5. Enter the name of the route map in the Route Map Name field. In this example, enter the name DAMPENING_MAP_ENABLE for the "empty" route map that you will use to globally enable dampening.
6. Enter the sequence (instance) number in the Sequence field or use the default value.
7. 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. In this example, select Permit.
8. Click the Add button to apply the changes to the device's running-config file.

NOTE: In this case, you are configuring an "empty" route map with no match or set statements, so you do not need to select the Route Map Match or Route Map Set link.

9. Enter the name of the route map you will use to set dampening parameters for a neighbor in the Route Map Name field. In this example, enter the name DAMPENING_MAP_NEIGHBOR_A.
10. 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. In this example, select Permit.
11. Click the Add button to apply the changes to the device's running-config file.
12. Select the Route Map Set link to display the BGP Route Map Set panel.

NOTE: If the interface displays a table listing the configured route maps, select the Route Map Set link under the table or click Modify next to the row describing the route map you are configuring.

13. Select the route map name and sequence from the Route Map Name.Sequence field's pulldown menu.
14. Select the checkbox in the Dampening section to enable dampening for routes that match the route map.
15. Click the Apply button to apply the changes to the device's running-config file.
16. In the tree view, under BGP in the Configure section, click on the General link to display the BGP configuration panel.
17. In the Dampening section, click next to Route-Map, then select the dampening route map from the Route-Map field's pulldown menu. In this example, select the map named DAMPENING_MAP_ENABLE.

NOTE: The route map appears in this menu only if you have already configured the route map.

18. Click Apply to apply the changes to the device's running-config file.
19. In the tree view, under BGP in the Configure section, click on the Neighbor link to display the list of BGP neighbors.
20. Select the Modify button to the right of the row describing the neighbor to which you want to apply the dampening route map you configured in steps 9 – 15.

21. Select the Route Map link at the bottom of the panel to display the BGP Neighbor Route Map panel, as shown in the following example.

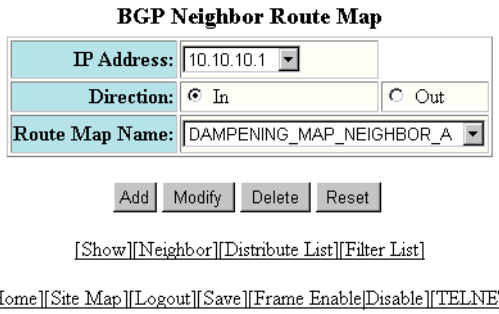

22. Select the neighbor IP address from the IP Address field's pulldown menu.
23. Select the traffic direction to which you want to apply the route map. You can select In or Out. In this example, select In.
24. Select the route map from the Route Map Name field's pulldown menu. In this example, select DAMPENING_MAP_NEIGHBOR_A.
25. Click Add to apply the changes to the device's running-config file.
26. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Removing Route Dampening from a Route

You can un-suppress routes by removing route flap dampening from the routes. The routing switch allows you to un-suppress all routes at once or un-suppress individual routes.

To un-suppress routes, use either of the following methods.

USING THE CLI

To un-suppress all the suppressed routes, enter the following command at the Privileged EXEC level of the CLI:

```
HP9300# clear ip bgp damping
```

Syntax: clear ip bgp damping [<ip-addr> <ip-mask>]

The <ip-addr> parameter specifies a particular network.

The <ip-mask> parameter specifies the network's mask.

To un-suppress a specific route, enter a command such as the following:

```
HP9300# clear ip bgp damping 209.157.22.0 255.255.255.0
```

This command un-suppresses only the route(s) for network 209.157.22.0/24.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Command in the tree view to expand the list of command options.
3. Click on the Clear link to display the Clear panel.
4. Select the checkbox next to BGP Dampening.

5. Specify the routes from which you want to remove dampening:
   - To clear dampening for all routes, select the All option.
   - To clear dampening for a specific route, select IP, then enter the network address and sub-net mask in the IP and Mask fields.
6. Click the Apply button to implement the change.

Displaying and Clearing Route Flap Dampening Statistics

The software provides many options for displaying and clearing route flap statistics. To display the statistics, use either of the following methods.

Displaying Route Flap Dampening Statistics

To display route flap dampening statistics, use the following CLI method.

USING THE CLI

To display route dampening statistics or all the dampened routes, enter the following command at any level of the CLI:

```
HP9300# show ip bgp flap-statistics

Total number of flapping routes: 414
    Status Code  >:best d:damped h:history *:valid
    Network            From            Flaps Since    Reuse    Path
h>  192.50.206.0/23    166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 701
h>  203.255.192.0/20   166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 7018
h>  203.252.165.0/24   166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 7018
h>  192.50.208.0/23    166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 701
h>  133.33.0.0/16      166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 701
*>  204.17.220.0/24    166.90.213.77   1     0 :1 :4  0 :0 :0  65001 4355 701 62
```

Syntax: show ip bgp flap-statistics [regexp <regular-expression> | <address> <mask> [longer-prefix] | neighbor <ip-addr>]

The **regexp** <regular-expression> parameter is a regular expression. The regular expressions are the same ones supported for BGP4 AS-path filters. See “Using Regular Expressions” on page 8-44.

The <address> <mask> parameter specifies a particular route. If you also use the optional **longer-prefix** parameter, then all statistics for routes that match the specified route or have a longer prefix than the specified route are displayed. For example, if you specify 209.157.0.0 longer, then all routes with the prefix 209.157. or that have a longer prefix (such as 209.157.22.) are displayed.

The **neighbor** <ip-addr> parameter displays route flap dampening statistics only for routes learned from the specified neighbor. You also can display route flap statistics for routes learned from a neighbor by entering the following command: **show ip bgp neighbor** <ip-addr> **flap-statistics**.

This display shows the following information.

Table 8.2: Route Flap Dampening Information

This Field...	Displays...
Total number of flapping routes	The total number of routes in the routing switch's BGP4 route table that have changed state and thus have been marked as flapping routes.
Status code	Indicates the dampening status of the route, which can be one of the following: • > – This is the best route among those in the BGP4 route table to the route's destination. • d – This route is currently dampened, and thus unusable. • h – The route has a history of flapping and is unreachable now. • * – The route has a history of flapping but is currently usable.
Network	The destination network of the route.
From	The neighbor that sent the route to the routing switch.
Flaps	The number of flaps (state changes) the route has experienced.
Since	The amount of time since the first flap of this route.
Reuse	The amount of time remaining until this route will be un-suppressed and thus be usable again.
Path	Shows the AS-path information for the route.

You also can display all the dampened routes by entering the following command:
show ip bgp dampened-paths.

USING THE WEB MANAGEMENT INTERFACE

You cannot display dampening statistics using the Web management interface.

Clearing Route Flap Dampening Statistics

To clear route flap dampening statistics, use the following CLI method.

NOTE: Clearing the dampening statistics for a route does not change the dampening status of the route.

USING THE CLI

To clear all the route dampening statistics, enter the following command at any level of the CLI:

```
HP9300# clear ip bgp flap-statistics
```

Syntax: `clear` ip bgp flap-statistics [regexp <regular-expression> | <address> <mask> | neighbor <ip-addr>]

The parameters are the same as those for the **show ip bgp flap-statistics** command (except the **longer** option is not supported). See "Displaying Route Flap Dampening Statistics" on page 8-74.

NOTE: The **clear ip bgp damping** command not only clears statistics but also un-suppresses the routes. See "Displaying Route Flap Dampening Statistics" on page 8-74.

USING THE WEB MANAGEMENT INTERFACE

You cannot clear dampening statistics using the Web management interface.

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
- Route flap dampening statistics

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 : 64512
  Confederation Identifier : 10
  Confederation Peers : 64512 64513
  Maximum Number of Attribute Entries Supported :10000
  Maximum Number of Routes Supported : 60000
  Maximum Number of Neighbors Supported : 3
  Maximum Number of Paths Supported for Load Sharing : 1
  Number of Routes Installed : 58756
  Number of Attribute Entries Installed : 7750

  Neighbor Address AS#    State  StateChangeTime  RtReceived  RtInstalled  RtSent
  192.168.11.1     64512  ESTAB  0  :0  :43 :54   65871       65871        0
  192.168.88.28    64512  ESTAB  0  :2  :26 :43   1           1            65875
  192.168.199.1    64513  ESTAB  0  :0  :48 :5    0           0            65875
```

This display shows the following information.

Table 8.3: BGP4 Summary Information

This Field...	Displays...
Local AS Number	The BGP4 AS number the router is in.
Confederation Identifier	The AS number of the confederation the routing switch is in.
Confederation Peers	The numbers of the local ASs contained in the confederation. This list matches the confederation peer list you configure on the routing switch.

Table 8.3: BGP4 Summary Information (Continued)

This Field...	Displays...
Maximum Number of Attribute Entries Supported	The number of attribute entries the router'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 8-23.
Maximum Number of Routes Supported	The number of BGP4 routes the router's memory can hold. To reconfigure the memory size for routes, see "Changing the Maximum Number of Routes" on page 8-22.
Maximum Number of Neighbors Supported	The number of BGP4 neighbors the router can have. To reconfigure the memory size for neighbors, see "Changing the Maximum Number of Neighbors" on page 8-21.
Maximum Number of Paths Supported for Load Sharing	The maximum number of route paths across which the device can balance traffic to the same destination. BGP4 load sharing is a new feature in software release 06.x. The feature is enabled by default but the default number of paths is 1. You can increase the number from 2 – 8 paths. See "BGP4 Load Sharing".
Number of Routes Installed	The number of BGP4 routes in the router's BGP4 route table. To display the BGP4 route table, see "Displaying the BGP4 Route Table" on page 8-87.
Number of Attribute Entries Installed	The number of BGP4 route-attribute entries in the router's route-attributes table. To display the route-attribute table, see "Displaying BGP4 Route-Attribute Entries" on page 8-90.
Neighbor Address	The IP addresses of this router's BGP4 neighbors.

Table 8.3: BGP4 Summary Information (Continued)

This Field...	Displays...
State	The state of this router's neighbor session with each neighbor. The states are from this router'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 router: • 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 router 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.
RtReceived	The total number or routes received in UPDATE messages from the neighbor since the session was first established.
RtInstalled	The number of routes received from the neighbor that this router actually installed in the BGP4 route table. Usually, this number is lower than the RoutesRcvd number. The difference indicates that this router filtered out some of the routes received in the UPDATE messages.
RtSent	The number of BGP4 routes that the routing switch has sent to the neighbor.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of monitoring options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Summary link to display the BGP Neighbor Summary panel.

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 neighbor
```

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                Confed_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
```

```
TCP Connection state: ESTABLISHED
Local host:  192.168.11.1, Local  Port: 8180
Remote host: 192.168.11.2, Remote Port: 179
ISentSeq:  710279168  SendNext:  710279383  TotUnAck:          0
SendWnd:       16384  TotSent:         215  ReTrans:         171
IRcvSeq:           0  RcvNext:         462  RcvWnd:        16384
TotalRcv:        462  RcvQue:            0  SendQue:           0
```

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.

The TCP statistics at the end of the display show status for the TCP session with the neighbor. Most of the fields show information stored in the routing switch's Transmission Control Block (TCB) for the TCP session between the routing switch and its neighbor. These fields are described in detail in section 3.2 of RFC 793, "Transmission Control Protocol Functional Specification".

This display shows the following information.

Table 8.4: BGP4 Neighbor Information

This Field...	Displays...
Total number of BGP Neighbors	The total number of neighbors configured on this router.
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, an EBGP session, or a confederation EBGP session. • EBGP – The neighbor is in another AS. • EBGP_Confed – The neighbor is a member of another sub-AS in the same confederation. • IBGP – The neighbor is in the same AS.

Table 8.4: BGP4 Neighbor Information (Continued)

This Field...	Displays...
State	The state of the router's session with the neighbor. The states are from this router'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 router: • 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 router 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 router sends keep alive messages to the neighbor. See "Changing the Keep Alive Time and Hold Time" on page 8-20.
Hold Time	The hold time, which specifies how many seconds the router 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 8-20.

Table 8.4: 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 8-14. **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 router 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 router 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 router.
Update	For the Message Sent and Message Received columns, indicates how many UPDATE messages have been sent and received by this router.
Notifications	For the Message Sent and Message Received columns, indicates how many NOTIFICATION messages have been sent and received by this router.
Open	For the Message Sent and Message Received columns, indicates how many OPEN messages have been sent and received by this router.

Table 8.4: 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 8.4: BGP4 Neighbor Information (Continued)

This Field...	Displays...

Notification Message Error Code Received	If the router 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 8.4: 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 router 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.
TCP Connection state	The state of the connection with the neighbor. The connection can have one of the following states: • LISTEN – Waiting for a connection request. • SYN-SENT – Waiting for a matching connection request after having sent a connection request. • SYN-RECEIVED – Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTABLISHED – Data can be sent and received over the connection. This is the normal operational state of the connection. • FIN-WAIT-1 – Waiting for a connection termination request from the remote TCP, or an acknowledgment of the connection termination request previously sent. • FIN-WAIT-2 – Waiting for a connection termination request from the remote TCP. • CLOSE-WAIT – Waiting for a connection termination request from the local user. • CLOSING – Waiting for a connection termination request acknowledgment from the remote TCP. • LAST-ACK – Waiting for an acknowledgment of the connection termination request previously sent to the remote TCP (which includes an acknowledgment of its connection termination request). • TIME-WAIT – Waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request. • CLOSED – There is no connection state.
Local host	The IP address of the routing switch.
Local port	The TCP port the routing switch is using for the BGP4 TCP session with the neighbor.
Remote host	The IP address of the neighbor.
Remote port	The TCP port the neighbor is using for the BGP4 TCP session with the routing switch.
ISentSeq	The initial send sequence number for the session.
SendNext	The next sequence number to be sent.

Table 8.4: BGP4 Neighbor Information (Continued)

This Field...	Displays...
TotUnAck	The number of sequence numbers sent by the routing switch that have not been acknowledged by the neighbor.
SendWnd	The size of the send window.
TotSent	The number of sequence numbers sent to the neighbor.
ReTrans	The number of sequence numbers that the routing switch retransmitted because they were not acknowledged.
IRcvSeq	The initial receive sequence number for the session.
RcvNext	The next sequence number expected from the neighbor.
RcvWnd	The size of the receive window.
TotalRcv	The number of sequence numbers received from the neighbor.
RcvQue	The number of sequence numbers in the receive queue.
SendQue	The number of sequence numbers in the send queue.

To display detailed neighbor information, enter a command such as the following:

```
HP9300A(config)# show ip bgp neighbor 192.168.12.3 routes-summary

Neighbor Address                          :192.168.12.3
Number of Routes Received                 :0
Number of Routes Accepted                 :0
Number of Routes Filtered                 :0
Number of NLRIS Received in UM            :0
Number of NLRIS Withdrawn Received in UM :0
Number of NLRIS Replace Received in UM    :0
Number of Routes Advertised               :65871
Number of Routes To Send                  :0
Number of NLRIS Sent in UM                :66632
Number of NLRIS Replaced in UM            :686
Number of NLRIS Withdrawn in UM           :75
Number of NLRIS to Withdraw               :0
```

This display shows the following information.

Table 8.5: Detailed Neighbor Information

Field	Description
Neighbor Address	The IP address of the neighbor for which you are displaying information.
Number of Routes Received	The total number of routes the routing switch has received from this neighbor.
Number of Routes Accepted	The number of routes the routing switch has accepted from the neighbor. This number is always equal to or less than the number of routes received. If the number is less, then the number of routes accepted plus the number of routes filtered equals the number of routes received.

Table 8.5: Detailed Neighbor Information (Continued)

Field	Description
Number of Routes Filtered	The number of routes received from the neighbor that the routing switch filtered out.
Number of NLRIs Received in UM	The number of routes received in Network Layer Reachability (NLRI) format in UPDATE messages (UMs).
Number of NLRIs Withdrawn Received in UM	The number of routes withdrawn by NLRI information received in UPDATE messages (UMs).
Number of NLRIs Replace Received in UM	The number of routes replaced by updated versions of the routes in NLRI information received in UPDATE messages (UMs).
Number of Routes Advertised	The number of routes the routing switch has advertised to this neighbor.
Number of Routes To Send	The number of routes the routing switch has queued to send to this neighbor.
Number of NLRIs Sent in UM	The number of NLRIs for new routes the routing switch has sent to this neighbor in UPDATE messages.
Number of NLRIs Replaced in UM	The number of NLRIs with replacement routes the routing switch has sent to this neighbor in UPDATE messages.
Number of NLRIs Withdrawn in UM	The number of NLRIs to withdraw routes the routing switch has sent to this neighbor in UPDATE messages.
Number of NLRIs to Withdraw	The number of NLRIs for withdrawing routes the routing switch has queued up to send to this neighbor in UPDATE messages.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of monitoring options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Neighbor link to display the BGP Neighbor Statistics panel.

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 8-3 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>:<num> | local-as | no-export | no-advertise | internet]
[community-list <num, 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 **local-as**, **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 8-44.

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...

You also can display only the routes that match a specified address and mask, by entering a command such as the following.

```
HP9300 Router# show ip bgp 3.3.0.0/16 longer
        Number of BGP Routes matching display condition : 2
Status codes: s suppressed, * valid, > best, i internal
Origin codes: i - IGP, e - EGP, ? - incomplete
    Network          Next Hop        Metric     LocPrf     Weight Path
*>  3.3.3.0          207.95.6.101    0          100        0      2 ?
*>  3.3.4.0          207.95.6.101    0          100        0      2 ?
```

This example shows all the routes for networks beginning with 3.3. The mask value and **longer** parameter specify the range of network addresses to be displayed. In this example, all routes within the range 3.3.0.0 – 3.3.255.255 are listed.

Syntax: show ip bgp <ip-addr> <ip-mask> longer-prefixes

or

Syntax: show ip bgp <ip-addr>/<mask-bits> longer-prefixes

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

Table 8.6: BGP4 Route Information

This Field...	Displays...
Total number of BGP Routes	The number of routes contained in this router'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 router associates with routes from a specific neighbor. For example, if the router receives routes to the same destination from two BGP4 neighbors, the router 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 8-3. • 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. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of monitoring options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Routes link to display the BGP Routes panel.

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 bgp attribute-entries
```

Syntax: show ip bgp 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 8.7: 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.

Table 8.7: BGP4 Route-Attribute Entries Information (Continued)

This Field...	Displays...
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. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of monitoring options.
3. Click on the plus sign next to BGP in the tree view to expand the list of BGP option links.
4. Click on the Attributes link to display the BGP Attributes 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. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of monitoring options.
3. Click on the plus sign next to IP in the tree view to expand the list of IP option links.
4. Click on the Routing Table link to display the IP route table.

Displaying Route Flap Dampening Statistics

To display route flap dampening statistics, use the following CLI method.

USING THE CLI

To display route dampening statistics or all the dampened routes, enter the following command at any level of the CLI:

```
HP9300# show ip bgp flap-statistics

Total number of flapping routes: 414
    Status Code  >:best d:damped h:history *:valid
    Network           From            Flaps Since    Reuse    Path
h>  192.50.206.0/23   166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 701
h>  203.255.192.0/20  166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 7018
h>  203.252.165.0/24  166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 7018
h>  192.50.208.0/23   166.90.213.77   1     0 :0 :13 0 :0 :0  65001 4355 1 701
```

```
h>  133.33.0.0/16       166.90.213.77   1      0 :0 :13 0 :0 :0  65001 4355 1 701
*>  204.17.220.0/24     166.90.213.77   1      0 :1 :4  0 :0 :0  65001 4355 701 62
```

Syntax: show ip bgp flap-statistics [regexp <regular-expression> | <address> <mask> [longer-prefix] | neighbor <ip-addr>]

The **regexp** <regular-expression> parameter is a regular expression. The regular expressions are the same ones supported for BGP4 AS-path filters. See “Using Regular Expressions” on page 8-44.

The <address> <mask> parameter specifies a particular route. If you also use the optional **longer-prefix** parameter, then all statistics for routes that match the specified route or have a longer prefix than the specified route are displayed. For example, if you specify 209.157.0.0 longer, then all routes with the prefix 209.157. or that have a longer prefix (such as 209.157.22.) are displayed.

The **neighbor** <ip-addr> parameter displays route flap dampening statistics only for routes learned from the specified neighbor. You also can display route flap statistics for routes learned from a neighbor by entering the following command: **show ip bgp neighbor** <ip-addr> **flap-statistics**.

This display shows the following information.

Table 8.8: Route Flap Dampening Information

This Field...	Displays...
Total number of flapping routes	The total number of routes in the routing switch’s BGP4 route table that have changed state and thus have been marked as flapping routes.
Status code	Indicates the dampening status of the route, which can be one of the following: • > – This is the best route among those in the BGP4 route table to the route’s destination. • d – This route is currently dampened, and thus unusable. • h – The route has a history of flapping and is unreachable now. • * – The route has a history of flapping but is currently usable.
Network	The destination network of the route.
From	The neighbor that sent the route to the routing switch.
Flaps	The number of flaps (state changes) the route has experienced.
Since	The amount of time since the first flap of this route.
Reuse	The amount of time remaining until this route will be un-suppressed and thus be usable again.
Path	Shows the AS-path information for the route.

You also can display all the dampened routes by entering the following command:
show ip bgp dampened-paths.

USING THE WEB MANAGEMENT INTERFACE

You cannot display dampening statistics using the Web management interface.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Command in the tree view to expand the list of command options.
3. Click on the Clear link to display the Clear panel.
4. 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.
5. Click the Apply button to implement the change.

Clearing Route Flap Dampening Statistics

To clear route flap dampening statistics, use the following CLI method.

NOTE: Clearing the dampening statistics for a route does not change the dampening status of the route.

USING THE CLI

To clear all the route dampening statistics, enter the following command at any level of the CLI:

```
HP9300# clear ip bgp flap-statistics
```

Syntax: clear ip bgp flap-statistics [regexp <regular-expression> | <address> <mask> | neighbor <ip-addr>]

The parameters are the same as those for the **show ip bgp flap-statistics** command (except the **longer** option is not supported). See “Displaying Route Flap Dampening Statistics” on page 8-74.

NOTE: The **clear ip bgp damping** command not only clears statistics but also un-suppresses the routes. See “Displaying Route Flap Dampening Statistics” on page 8-74.

USING THE WEB MANAGEMENT INTERFACE

You cannot clear dampening statistics using the Web management interface.

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 HP 9304M, HP 9308M, and HP 6208M-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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Command in the tree view to expand the list of command options.
3. Click on the Clear link to display the Clear panel.
4. Select BGP Neighbor Soft-Outbound.
5. Use the default value All to resend the BGP4 route table to all neighbors or select a neighbor from the field's pulldown menu.
6. Click the Apply button to implement the change.

Removing Route Flap Dampening

You can un-suppress routes by removing route flap dampening from the routes. The routing switch allows you to un-suppress all routes at once or un-suppress individual routes.

To un-suppress routes, use either of the following methods.

USING THE CLI

To un-suppress all the suppressed routes, enter the following command at the Privileged EXEC level of the CLI:

```
HP9300# clear ip bgp damping
```

Syntax: clear ip bgp damping [<ip-addr> <ip-mask>]

The <ip-addr> parameter specifies a particular network.

The <ip-mask> parameter specifies the network's mask.

To un-suppress a specific route, enter a command such as the following:

```
HP9300# clear ip bgp damping 209.157.22.0 255.255.255.0
```

This command un-suppresses only the route(s) for network 209.157.22.0/24.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Command in the tree view to expand the list of command options.
3. Click on the Clear link to display the Clear panel.
4. Select the checkbox next to BGP Dampening.
5. Specify the routes from which you want to remove dampening:
   - To clear dampening for all routes, select the All option.
   - To clear dampening for a specific route, select IP, then enter the network address and sub-net mask in the IP and Mask fields.
6. Click the Apply button to implement the change.

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 8-79.

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. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Command in the tree view to expand the list of command options.
3. Click on the Clear link to display the Clear panel.
4. 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.
5. Click the Apply button to implement the change.