
Chapter 14

BGP4 Commands

address-filter

Configures an address filter for filtering routes in BGP4 updates based on IP address.

EXAMPLE:

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 <ip-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> <ip-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.

Possible values: see above

Default value: N/A

aggregate-address

Configures the routing switch to aggregate routes in a range of networks into a single CIDR number.

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.

EXAMPLE:

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> <ip-mask> [as-set] [summary-only] [suppress-map <map-name>] [advertise-map <map-name>] [attribute-map <map-name>]

The <ip-addr> and <ip-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 the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide* for information on defining a route map.

Possible values: see above

Default value: N/A

always-compare-med

Configures the routing switch to compare the MEDs for all paths for a route, regardless of the AS through which the paths pass. By default, the routing switch compares the MED values only among paths through the same AS.

EXAMPLE:

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

Possible values: N/A

Default value: Disabled

as-path-filter

Configures an AS-path filter for filtering routes in BGP4 updates based on AS-path.

EXAMPLE:

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 HP 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.

NOTE: You can use regular expressions as part of the AS-path. See the "Configuring BGP4" chapter of the *Advanced Configuration and Management Guide*.

Possible values: N/A

Default value: Disabled

auto-summary

Enables or disables auto summary. The auto summary feature summarizes the routes it redistributes from IGP to BGP4. The routing switch summarizes sub-nets into their natural class A, B, or C networks. For example, if an AS contains sub-nets 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 sub-nets in its advertisements to BGP4 neighbors as 1.0.0.0/8. The auto summary feature is disabled by default.

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 the "Configuring BGP4" chapter of the *Advanced Configuration and Management Guide*.

EXAMPLE:

To enable auto summary, enter the following command:

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

Syntax: [no] auto-summary

Possible values: N/A

Default value: Disabled

client-to-client-reflection

Disables or re-enables route reflection. For more information about route reflection, see the "Configuring BGP4" chapter of the *Advanced Configuration and Management Guide*.

EXAMPLE:

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

Possible values: N/A

Default value: Enabled

cluster-id

Changes the BGP4 cluster ID. Use this command only on a BGP4 routing switch that you are using as a route reflector. For more information about route reflection, see the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide*.

EXAMPLE:

Enter the following commands to configure an HP routing switch as a route reflector for two neighbors in cluster 1:

```
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>

Possible values: 1 – 4294967295

Default value: the router ID, expressed as a 32-bit number

community-filter

Configures a community address filter for filtering routes in BGP4 updates based on community.

EXAMPLE:

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 routing switch applies the filters in numerical order, beginning with the lowest-numbered filter. When a filter match is true, the routing switch stops and does not continue applying filters from the list.

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

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

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

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

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

The **local-as** parameter specifies a community value. If a BGP4 route received by the routing switch from a peer has the community type LOCAL_AS, the routing switch advertises the route only within the sub-AS. If the route has the community type NO_EXPORT, the routing switch advertises the route only within the confederation.

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.

Possible values: N/A

Default value: Disabled

confederation

Configures a routing switch to be a member of a BGP confederation.

EXAMPLE:

To configure a routing switch to be a member of confederation 10, consisting of two sub-ASs (64512 and 64513):

```
HP9300A(config-bgp-router)# confederation identifier 10
HP9300A(config-bgp-router)# confederation peers 64512 64513
```

Syntax: confederation identifier <num>

Syntax: confederation peers <num> [<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.

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.

Possible values: 1 – 65535

Default value: N/A

dampening

Configures route flap dampening.

This section shows how to globally configure dampening. You also can use route maps to configure dampening for specific neighbors and routes.

EXAMPLE:

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.

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.

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.

Possible values: See above

Default value: Disabled

default-information-originate

Enables the routing switch to advertise a default BGP4 route.

EXAMPLE:

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

Possible values: N/A

Default value: Enabled

default-local-preference

Changes the local preference. The 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.

EXAMPLE:

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>

Possible values: 0 – 4294967295

Default value: 100

default-metric

Sets the default BGP4 MED (metric), 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.

EXAMPLE:

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

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

Syntax: default-metric <num>

Possible values: 0 – 4294967295

Default value: 0

distance

Changes the administrative distance for IBGP, EBGP, or Local BGP routes. 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.

See the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide* for a list of the default administrative distances for all types of routes.

EXAMPLE:

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 default is 20.

The <internal-distance> sets the IBGP distance and can be a value from 1 – 255. The default is 200.

The <local-distance> sets the Local BGP distance and can be a value from 1 – 255. The default is 200.

Possible values: see above

Default value: see above

end

Moves activity to the privileged EXEC level from any level of the CLI except the user EXEC level.

EXAMPLE:

To move to the privileged level, enter the following from any level of the CLI.

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

```
HP9300#
```

Syntax: end

Possible values: N/A

Default value: N/A

exit

Moves activity up one level from the current level. In this case, activity will be moved to the global level.

EXAMPLE:

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

```
HP9300(config)#
```

Syntax: exit

Possible values: N/A

Default value: N/A

fast-external-fallover

Configures the routing switch to immediately close the BGP4 session and TCP connection to locally attached neighbors that die. This feature is disabled by default.

EXAMPLE:

To enable fast external fallover, enter the following command:

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

Syntax: [no] fast-external-fallover

Possible values: N/A

Default value: Disabled

local-as

The local AS number identifies the AS the HP routing switch is in. The AS number can be from 1 – 65535. AS numbers 64512 – 65535 are the well-known private BGP4 AS numbers. There is no default local AS number.

EXAMPLE:

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

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

Syntax: local-as <num>

Possible values: 1 – 65535

Default value: None

max-attribute-entries

Changes the maximum number of BGP4 route attribute entries the routing switch can have.

EXAMPLE:

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>

Possible values:

See the “Memory Considerations” section of the “Configuring BGP” chapter in the *Advanced Configuration and Management Guide* for a listing of the maximum number of BGP4 route attribute entries the routing switch can have.

maximum-paths

Changes the maximum number of shared 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.

EXAMPLE:

To change the maximum number of shared paths:

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

Syntax: [no] maximum-paths <num>

Possible values: 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.

Default value: The default is 1.

max-neighbors

Changes the maximum number of BGP4 neighbors the routing switch can have.

EXAMPLE:

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

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

```
HP9300# reload
```

Syntax: max-neighbors <num>

Possible values: See the “Memory Considerations” section of the “Configuring BGP” chapter in the *Advanced Configuration and Management Guide* for a listing of the maximum number of BGP4 route attribute entries the routing switch can have.

max-routes

Changes the maximum number of BGP4 routes the routing switch can have.

EXAMPLE:

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>

Possible values: See the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide* for a listing of the maximum number of BGP4 routes the routing switch can have.

neighbor

Adds a BGP4 neighbor (peer). In addition to identifying the neighbor’s IP address and AS number, you can set other parameters that control the routing switch’s interaction with the neighbor.

EXAMPLE:

You can add a neighbor by specifying just the IP address and AS number. To set additional options, see the syntax descriptions below.

```
HP9300(config-bgp-router)# neighbor 1.1.1.0 remote-as 1
```

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 routing switch applies

the filters in the order in which you list them and stops applying the filters in the distribute list when a match is found.

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

NOTE: The address filter must already be configured.

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 routing switch applies the filters in the order in which you list them and stops applying the filters in the AS-path filter list when a match is found. 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 “ip as-path” on page 6-23.

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 routing switch and Chassis device 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 the “Changing the Maximum Number of Routes” section in the *Advanced Configuration and Management Guide*.

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

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 “ip prefix-list” on page 6-30.

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

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

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

NOTE: The route map must already be configured.

route-reflector-client specifies that this neighbor is a route-reflector client of the routing switch. Use the parameter only if this routing switch is going to be a route reflector. This option is disabled by default.

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

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 routing switch 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.

update-source loopback <num> configures the routing switch to communicate with the neighbor through the loopback address on the specified interface. Using a loopback address for neighbor communication avoids problems that can be caused by unstable 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.

Possible values: see above

Default value: N/A

network

Specifies a list of networks for the routing switch to advertise through BGP4.

EXAMPLE:

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.

Possible values: see above

Default value: N/A

no

Disables other commands. To disable a command, place the word **no** before the command.

quit

Returns you from any level of the CLI to the User EXEC mode.

EXAMPLE:

```
HP9300(config-bgp-router)# quit
HP9300>
```

Syntax: quit

Possible values: N/A

Default value: N/A

redistribute connected

Configures parameters for redistributing routes to directly attached devices into BGP4. Redistribution into BGP4 is disabled by default.

EXAMPLE:

To configure the routing switch to redistribute routes to directly attached devices, enter the following command:

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

Syntax: redistribute connected [metric <num>] [route-map <map-name>] [weight <num>]

The **connected** parameter indicates that you are redistributing routes to directly attached devices 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 the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide* 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.

Possible values: see above

Default value: see above

redistribute ospf

Configures parameters for redistributing OSPF routes into BGP4. Redistribution into BGP4 is disabled by default.

EXAMPLE:

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 the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide* 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.

Possible values: see above

Default value: see above

redistribute rip

Configures parameters for redistributing IP/RIP routes into BGP4. Redistribution into BGP4 is disabled by default.

EXAMPLE:

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 the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide* 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.

Possible values: see above

Default value: see above

redistribute static

Configures parameters for redistributing static routes into BGP4. Redistribution into BGP4 is disabled by default.

EXAMPLE:

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 the “Configuring BGP4” chapter of the *Advanced Configuration and Management Guide* 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.

Possible values: see above

Default value: see above

show

Displays a variety of configuration and statistical information about the switch or routing switch. See “Show Commands” on page 20-1.

synchronization

Enables or disables synchronization. When synchronization is enabled, the routing switch waits until the IGP's in the local AS have fully exchanged route information before BGP4 advertises the routes to its remote BGP4 neighbors.

EXAMPLE:

To enable synchronization, enter the following command:

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

Syntax: [no] synchronization

Possible values: N/A

Default value: Disabled

table-map

Configures an existing route map to change the route tag in routes when adding them to the IP route table.

EXAMPLE:

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
```

Syntax: table-map <route-map>

Possible values: a route-map name

Default value: N/A

timers

Sets the BGP4 Keep Alive Time and Hold Time on the routing switch.

EXAMPLE:

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>

Possible values:

Keep Alive Time 0 – 65535.

Hold Time 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.

Default value:

Keep Alive Time 60 seconds.

Hold Time 180 seconds.

write memory

Saves the running configuration into the startup-config file.

EXAMPLE:

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

Syntax: write memory

Possible values: N/A

Default value: N/A

write terminal

Displays the running configuration of the HP switch or routing switch on the terminal screen.

NOTE: This command is equivalent to the **show running-config** command.

EXAMPLE:

```
HP9300(config-bgp-router)# wr term
```

Syntax: write terminal

Possible values: N/A

Default value: N/A

