

12 Configuring BGP

BGP (Border Gateway Protocol) is a protocol for exchanging routing information between gateway hosts in a network of AS's (autonomous systems). BGP is the most common protocol used between gateway hosts on the Internet. The routing table exchanged contains a list of known routers, the addresses they can reach, and a cost metric associated with the path to each router so that the best available route is chosen. The OmniCore implementation supports BGP-4, the latest version of BGP, as defined in RFC 1771.

Hosts using BGP communicate using the TCP (Transmission Control Protocol) on port 179. On connection start, BGP peers exchange complete copies of their routing tables, which can be quite large. However, only changes are exchanged after startup, which makes long running BGP sessions more efficient than shorter ones. BGP-4 lets administrators configure cost metrics based on policy statements.

BGP communicates with autonomous (local) networks using Internal BGP (IBGP). The routers inside the autonomous network thus maintain two routing tables: one for the interior gateway protocol and one for IBGP.

BGP-4 makes it easy to use Classless Inter-Domain Routing (CIDR), which is a way to have more addresses within the network than with the current Internet Protocol address assignment scheme. BGP's basic unit of routing information is the BGP path, which is a route to a certain set of CIDR prefixes. Paths are tagged with various path attributes, of which the most important are AS_PATH and NEXT_HOP.

One of BGP-4's most important functions is loop detection at the autonomous system level, using the AS_PATH attribute, a list of autonomous systems being used for data transport. The syntax of this attribute is made more complex by its need to support path aggregation, when multiple paths are collapsed into one to simplify further route advertisements. A simplified view of AS_PATH is that it is the list of Autonomous Systems that a route goes through to reach its destination. Loops are detected and avoided by checking for your own AS number in AS_PATH's received from neighboring Autonomous Systems.

◆ BGP 256 MB Upgrade Kit ◆

Implementing BGP on the OmniCore 5052 and 5022 requires an upgrade to 256 MB on the primary EMM and secondary EMM, if present. Contact your Alcatel sales representative about obtaining a BGP 256 MB Upgrade Kit for each EMM module in your system.

◆ BGP Restrictions on Use ◆

BGP on the OmniCore 5052 and 5022 can support up to 1.5 million paths if you have installed interface modules with the large table memory feature (recommended). Exceeding this number may result in degraded system performance. See [Appendix C, "Technical Specifications"](#) to determine whether a particular interface module has large table (LT) or small table (ST) capability.

The following BGP functionality is discussed in this chapter.

- Enabling BGP
- BGP Aggregation
- BGP Dampening
- BGP Networks
- BGP Paths
- BGP Policy
- BGP Redistribution Filters

BGP Commands

The BGP commands in the OmniCore CLI are listed in the following table and in additional tables later in this chapter. All BGP commands appear under the *ip bgp* location in the CLI tree. To see more information regarding BGP commands, see the *OmniCore CLI Reference Manual*.

BGP Commands

Command	Default	Description
ip bgp as	no default	Configures the autonomous system number for the BGP speaker.
ip bgp aspath-compare	enable	Enables or disables AS path comparison.
ip bgp cluster-id	0.0.0.0	Configures a BGP cluster ID when there are multiple route-reflectors in a cluster.
ip bgp dfault-lpref	100	Sets the default local preference value.
ip bgp fefo	disable	Enables or disables Fast External Fall Over.
ip bgp log-event	no default	Sets the BGP event log to display specified types of events.
ip bgp log-status	disable	Enables or disables BGP logging status.
ip bgp med-always	disable	Enables or disables Med comparison between different autonomous systems.
ip bgp missing-med	worst	Sets the missing MED parameter.
ip bgp route show	no default	Displays BGP route information.
ip bgp router-id	0.0.0.0	Configures the BGP router ID.
ip bgp rr	disable	Enables or disables the local BGP speaker to be a route-reflector.
ip bgp rtgen-interval	15 seconds	Sets a minimum AS origination interval in seconds.
ip bgp stats	none	Displays current BGP statistics.
ip bgp status	disable	Administratively enables or disables Border Gateway Protocol (BGP)
ip bgp sync	disable	Enables or disables IGP synchronization.
ip bgp tag	zero	Sets a BGP manual tag value.

Enabling BGP

Two commands are required to enable BGP routing on an OmniCore routing switch. You must issue the *ip bgp as* command to assign the AS number to the router and then enable BGP routing using the *ip bgp status* command, as shown in the example below. BGP routing is disabled by default.

```
OmniCore> ip bgp as 100
OmniCore> ip bgp status enable
```

You can view current BGP settings by issuing the *ip bgp show* command.

```
OmniCore> ip bgp show
Admin Status           :enable
Oper Status           :down
Autonomous System Number :200
BGP Router Id         :0.0.0.0
BGP Cluster Id        :0.0.0.0
IGP Synchronization Status :disable
Minimum as Origination Interval :15
Default Local Preference :100
Route Reflection      :disable
Missing Med status    :worst
Aspath Comparison     :enable
MED Compare Between different AS :disable
Fast External Fall Over :disable
Manual Tag            :0
Logging Status        :disable
```

Once BGP is enabled, you need to define the BGP peers that the local router will speak to. See [Configuring BGP Peers](#) on page 12-11.

BGP Aggregation

BGP aggregates are used to reduce the size of routing tables by combining the attributes of several different routes and allowing a single aggregate route to be advertised. Keep in mind that you cannot aggregate an address (for example, 100.10.0.0) if you do not have at least one more-specific route of the address (for example, 100.10.20.0) in the BGP routing table. Also note that only the aggregate is advertised unless aggregate summarization is disabled using the *ip bgp aggregate summarize* command.

You can view current BGP aggregates by issuing the *ip bgp aggregate show* command.

```
OmniCore> ip bgp aggregate show
Address      Mask          Summarize  As-Set   State    RowStatus
-----
10.10.0.0    255.255.255.0  enable    disable  inactive enable
-----
11.11.0.0    255.255.255.0  enable    disable  inactive enable
-----
Number of Entries Displayed: 2
```

The BGP aggregate commands are summarized below.

BGP Aggregate Commands

Command	Default	Description
ip bgp aggregate	no default	Creates, deletes, enables, or disables an aggregate route.
ip bgp aggregate as-set	disable	Enables or disables AS path aggregation.
ip bgp aggregate community	none	Sets the community string attribute for a BGP aggregate.
ip bgp aggregate lpref	0	Sets the local preference attribute value for a BGP aggregate.
ip bgp aggregate metric	0	Sets the MED attribute value for a BGP aggregate.
ip bgp aggregate summarize	enable	Enables or disables aggregate summarization.

BGP Dampening

BGP route dampening is a feature used to lessen the instability to a network caused by route flapping. Route flapping is regulated by configuring various dampening limits and duration parameters. A flap occurs each time a poorly behaved route is withdrawn from the table. Each flap adds 1 to the instability merit for the route. When the instability merit reaches the cutoff value (default: 3.0), the route is suppressed and is no longer advertised. The instability merit cannot rise above the ceiling value (default: 16.0). When a flapping route has stabilized, its instability merit is reduced by half for each half-life time duration. You can set separate half-life durations for reachable routes (default: 300 seconds) and unreachable routes (default: 900 seconds). When the instability merit falls below the reuse value (default: 2.0), the route is unsuppressed and is advertised once again.

In the OmniCore implementation, the dampening limits are defined as decimal values in the range 0 to 99.99. The OmniCore CLI will not allow you to set the reuse value greater than the cutoff value or the ceiling value less than the cutoff value. The duration values are defined in seconds. BGP route dampening is disabled by default.

You can view current BGP dampening settings by issuing the *ip bgp damp show* command.

```
OmniCore> ip bgp damp show
Dampening Status           :disable
Dampening Cutoff           :3.0
Dampening Reuse            :2.0
Dampening Ceiling          :16.0
Half Life while reachable  :300 secs
Half Life while unreachable :900 secs
Maximum instability time    :1800 secs
```

The various BGP dampening commands are summarized below.

BGP Dampening Commands

Command	Default	Description
ip bgp damp status	no default	Enables or disables BGP dampening.
ip bgp damp cutoff	3.0	Cutoff value at which a flapping route is suppressed.
ip bgp damp reuse	2.0	Reuse value when a suppressed route is unsuppressed.
ip bgp damp ceiling	16.0	Maximum values for instability merit.
ip bgp damp reach half-life	300 seconds	Half-life duration for a reachable route.
ip bgp damp unreach half-life	900 seconds	Half-life duration for an unreachable route.
ip bgp damp max-holdtime	1800 seconds	Maximum duration a route can be suppressed.
ip bgp damp clear	no default	Clears dampening history data.

BGP Networks

BGP network commands are used to create and enable network entries. The specified network must be known to the router, whether it is connected, static, or dynamically learned. Network entries indicate to BGP that this network should originate from this router. In the OmniCore implementation, BGP can accept up to 200 network entries. Individual commands allow you to set specific community, lpref, and metric values for each route generated by the *network* command.

You can view current BGP networks by issuing the *ip bgp network show* command.

```
OmniCore> ip bgp network show
Network      Mask          State      RowStatus
-----
1.10.10.10   255.255.255.255 active     enable
-----
1.10.65.0    255.255.255.252 active     enable
-----
1.10.201.0   255.255.255.252 active     enable
-----
Number of Entries Displayed: 3
```

The BGP network commands are summarized below.

BGP Network Commands

Command	Default	Description
ip bgp network	no default	Creates, deletes, enables, or disables a network.
ip bgp network community	"none"	Sets the community string attribute for the route generated by the network command.
ip bgp network lpref	0	Sets the local preference attribute value for the route generated by the network command.
ip bgp network metric	0	Sets the MED attribute value for the route generated by the network command.

BGP Paths

The `ip bgp path show` command displays information on every path currently in the path table. Since the number of paths may run into the thousands, this command provides a number of methods for displaying a specific path or matching entries for a portion of a path, peer address, and protocol type. Matching paths are displayed in table format, as shown in the example below.

```
OmniCore> ip bgp path 1.3.48.0 show
-----
Prefix/len      Peer Addr      Proto  Pref      State      Origin      NextHop
-----
1.3.48.0/30    1.11.48.2     ebgp   100       best       igp         1.11.48.2

Weight
AS Path: 64048

-----
1.3.48.0/30    1.11.51.2     ebgp   100       feasible   igp         1.11.51.2

Weight
AS Path: 64051 64003

-----
Number of Entries Displayed: 2
```

This path command allows you to filter the path list a table display based on a number of matching criteria. The example below shows the use of this command to display only the best paths.

```
OmniCore> ip bgp path best show
-----
Prefix/len      Peer Addr      Proto  Pref      State      Origin      NextHop
-----
1.3.48.0/30    1.11.65.2     ebgp   100       best       igp         1.11.65.2

Weight
AS Path: 64065 64003

-----
1.3.48.0/30    1.11.65.2     ebgp   100       best       igp         1.11.65.2

Weight
AS Path: 64065 64003
```

The example below shows the use of this command to display only paths matching the regular expression primitive (10 | 11). See [Regular Expressions](#) on page 12-8 for information on the use of regular expression in BGP commands.

```
OmniCore> path regexp (10 | 11) show
-----
Prefix/len      Peer Addr      Proto  Pref      State      Origin      NextHop
-----
205.222.0.0/16 1.10.65.2     ebgp   100       best       igp         1.10.65.2

AS Path: 64065 64777 3513 1239 1 3527 10

-----
205.222.0.0/16 1.10.201.2    ebgp   100       feasible   igp         1.10.201.2

AS Path: 64000 64017 64065 64777 3513 1239 1 3527 10

-----
Path Address                                :255.255.255.255
```

The BGP path commands are summarized below.

BGP Path Commands

Command	Default	Description
ip bgp path show	all paths	Displays information on a specific BGP path or multiple matching paths.
ip bgp path show table	all paths	Displays information on BGP paths that match the criteria of various filters.

Regular Expressions

Regular expressions are used to identify AS paths (such as are sent in certain BGP update messages to inform the receiving router of the route to a given network or set of networks) for purposes of making routing decisions. In this context, an AS path is a list of one or more unsigned 16-bit AS numbers, in the range 1 through 65535.

An ordinary pattern match string looks like "100 200" (without the double-quotes), which matches any AS path which contains the Autonomous System number 100 followed immediately by 200, anywhere within the AS path list. It would not match an AS path which was missing either number, or where the numbers didn't occur in the correct order, or where the numbers were not adjacent to one another.

Special pattern matching characters (sometimes called metacharacters) add the ability to specify that part of the pattern must match the beginning or end of the AS path list, or that some arbitrary number of AS numbers should match, etc. The following table defines the metacharacters used in the OmniCore BGP implementation.

Regular Expression Symbol Definitions

Symbol	Description
^	Matches the beginning of the AS path list.
123	Matches the AS number 123.
.	Matches any single AS number.
?	Matches zero or one occurrence of the previous token, which must be an AS number, a dot, an alternation, or a range.
+	Matches one or more occurrences of the previous token, which must be an AS number, a dot, an alternation, or a range.
*	Matches zero or more occurrences of the previous token, which must be an AS number, a dot, an alternation, or a range.
(Begins an alternation sequence of AS numbers. It matches any AS number listed in the alternation sequence.
	Separates AS numbers in an alternation sequence.
)	Ends an alternation sequence of AS numbers.
[Begin a range pair consisting of two AS numbers separated by a dash. It matches any AS number within that inclusive range.
-	Separates the endpoints of a range.
]	Ends a range pair.
\$	Matches the end of the AS path list.
, _	Commas, underscores (_), and spaces are ignored.

The regular expressions configured in the router are compared against an incoming AS path list one at a time until a match is found, or until all patterns have been unsuccessfully matched. Unlike some implementations, which use a character-based pattern matching logic, the OmniCore BGP implementation treats AS numbers as single tokens which provides two benefits:

- It makes writing (and reading) policies much easier.
- It enables the switch to begin using the policies more quickly after startup.

For example, to identify routes originating from internal autonomous systems, you would use the pattern:

```
[64512-65535]$
```

which means “match any AS number from 64512 to 65535 (inclusive) which occurs at the end of the AS path.” To accomplish the same thing using character-based pattern matching, you would have to use the following pattern:

```
(_6451[2-9]_|_645[2-9][0-9]_|_64[6-9][0-9][0-9]_|_65[0-9][0-9][0-9]_)$
```

Some examples of valid OmniCore regular expressions are shown in the following table.

Valid Regular Expression Examples

Example		Description
100	Meaning:	Any route which passes through AS number 100.
	Matches:	100 200 300 300 100 100
	Doesn't match:	200 300
^100	Meaning:	Any routes for which the next hop is AS number 100.
	Matches:	100 200 100
	Doesn't match:	50 100 200
100\$	Meaning:	Any route which originated from AS number 100 (AS numbers are prepended to the AS path list as they are passed on, so the originating AS is always the last number in the list).
	Matches:	100 200 200 100
	Doesn't match:	100 200
^100 500\$	Meaning:	A route with just two hops, 100 and 500.
	Matches:	100 500
	Doesn't match:	100 500 600 100 200 500
100 . . 200	Meaning:	Any routes with at least 4 hops, with 100 separated by any two hops from 200.
	Matches:	50 100 400 500 200 600 100 100 100 200
	Doesn't match:	100 200 100 100 200

Valid Regular Expression Examples (Continued)

(100 200) .+[500-650]\$	Meaning:	Any route which begins with 100 or 200, ends with an AS number between 500 and 650 (inclusive), and is at least three hops in length. The "." part matches at least one (but possibly more) AS numbers.
	Matches:	100 350 501 200 250 260 270 280 600
	Doesn't match:	100 600 100 400 600 700
^500\$	Meaning:	Only routes consisting of a single 500.
	Matches:	500
	Doesn't match:	500 500 100 500 600
[100-199]* 500 (900 950)\$	Meaning:	Any route which ends with any number of occurrences of AS numbers in the range 100- 199, followed by a 500, followed by either a 900 or a 950.
	Matches:	100 150 175 500 900 900 500 950
	Doesn't match:	100 200 500 900 100 199 500

Some examples of invalid OmniCore regular expressions are shown in the following table.

Invalid Regular Expression Examples

Error	Description
66543	Number is too large. AS numbers must be in the range 1 through 65535, inclusive.
64,512	Possibly an error, if the user meant the number "64512". The comma gets interpreted as a separator, thus this pattern is equivalent to the two AS numbers 64 and 512.
(100 200 300)	Alternation sequences must consist of single AS numbers separated by vertical bars, enclosed in parentheses.
(100* 200)	No metacharacters other than vertical bars may be included within an alternation sequence.
(100 (200 300))	Parentheses may not be nested. This pattern is actually equivalent to (100 200 300).
100 ^ 200	The '^' metacharacter must occur first in the pattern, as it matches the beginning of the AS path.
^500 \$ 600	The '\$' metacharacter must occur last in the pattern, as it matches the end of the AS path.
^? 100	The repetition metacharacters (?,+,*) cannot be applied to the beginning of the line. If it were legal, this pattern would be equivalent to the simpler pattern: 100
[1-(8 9)]*	A range cannot contain an alternation sequence.

Configuring BGP Peers

Two routers that have formed a TCP connection in order to exchange BGP routing data are called BGP “peers” or “neighbors.” BGP supports two types of neighbors, internal and external. For the purposes of this document, BGP running between routers in the same autonomous system are termed “internal” peers. BGP running between routers in different autonomous systems are termed “external” peers.

Two command are required to define a BGP peer. You must issue the *ip bgp peer* command to create the peer and then the *ip bgp peer remote-as* command to assign a remote AS number to the peer.

```
Omnicores> ip bgp peer 129.213.1.2 create
Omnicores> ip bgp peer 129.213.1.2 remote-as 200
```

You must enable BGP and issue similar peer commands on the peer router in order for the two routers to establish a TCP connection and begin exchanging BGP information.

You can view current BGP peers by issuing the *ip bgp peer show* command.

```
Omnicores> ip bgp peer show
Address      As      Type      Nbr Router ID  Admin  FSM
-----
1.10.65.2    64065   external  1.65.65.65    enable established
1.10.201.2   64000   external  1.201.201.201 enable  established
Number of Entries Displayed: 2
```

The BGP peer commands are summarized below.

BGP Peer Commands


Command	Default	Description
ip bgp peer	no default	Creates, deletes, administratively enables, or administratively disables a BGP peer.
ip bgp peer adv-interval	30 seconds	Configures the advertisement time interval between updates (in seconds) for a neighbor.
ip bgp peer auto-restart	enable	Enables or disables BGP peer automatic restart status.
ip bgp peer clear	no default	Restarts a peer.
ip bgp peer client	disable	Enables or disables peer route reflection client behavior.
ip bgp peer connect-interval	120 seconds	Configures the time interval between connection retries for a neighbor.
ip bgp peer default-originate	disable	Enables or disables BGP peer default origination.
ip bgp peer hold-interval	180 seconds	Configures the hold time interval between updates (in seconds) for a neighbor.
ip bgp peer in-aspathlist	“none”	Assigns an input AS path list filter list to the BGP peer.
ip bgp peer in-communitylist	“none”	Assigns an input community filter list to the BGP peer.
ip bgp peer in-map	“none”	Assigns an inbound policy for the BGP peer.

BGP Peer Commands (Continued)

ip bgp peer in-prefixlist	"none"	Assigns an input prefix filter list to the BGP peer.
ip bgp peer in-reconfigure	no default	Invokes inbound policy reconfiguration for the BGP peer.
ip bgp peer keep-alive	60 seconds	Configures the time interval between keep-alive messages for a neighbor.
ip bgp peer log	disable	Enables or disables BGP peer logging.
ip bgp peer max-prefix	4294967295	Sets the maximum number of prefixes that can be received from a peer.
ip bgp peer multi-hop	disable	Enables or disables BGP external peer multi-hop behavior.
ip bgp peer name	peer(<ipaddr>)	Sets the BGP peer name.
ip bgp peer nexthop-self	disable	Enables or disables BGP peer next hop set to self.
ip bgp peer out-aspathlist	"none"	Assigns an output AS path list filter list to the BGP peer.
ip bgp peer out-communitylist	"none"	Assigns an output community filter list to the BGP peer.
ip bgp peer out-map	"none"	Assigns an outbound policy for the BGP peer.
ip bgp peer out-prefixlist	"none"	Assigns an output prefix filter list to the BGP peer.
ip bgp peer out-reconfigure	no default	Invokes outbound policy reconfiguration for the BGP peer.
ip bgp peer passive	disable	Enables or disables BGP peer passive status.
ip bgp peer remote-as	zero	Assigns an autonomous system number to a BGP peer.
ip bgp peer remove-privateas	disable	Enables or disables BGP peer removal of private autonomous system numbers.
ip bgp peer soft-reconfig	enable	Enables or disables BGP peer soft reconfiguration on policy changes status.
ip bgp peer stats-clear	no default	Clears the statistics for a neighbor.

Configuring Internal and External Peers

The following diagram shows a configuration of internal and external BGP speakers. Routers B and C belong to the same autonomous system, AS200, and communicate using an IBGP (Interior BGP) protocol. Routers A and D exist in AS100 and AS300 respectively, and communicate with Routers B and C using EBGP (Exterior BGP) protocol.



Configuring BGP Internal and External Peers

For purposes of illustration, our example will configure the internal and external BGP peers shown above, ignoring specifics of IBGP configuration.

1. On **Router A** assign an AS number of 100, create an external peer of Router B in AS200, and enable BGP protocol.

```
Router A> ip bgp as 100
Router A> ip bgp peer 129.213.1.2 create
Router A> ip bgp peer 129.213.1.2 remote-as 200
Router A> ip bgp status enable
```

2. On **Router B** assign an AS number of 200, create an external peer of Router A in AS100, create an internal peer of Router C in AS200, and enable BGP protocol.

```
Router B> ip bgp as 200
Router B> ip bgp peer 129.213.1.1 create
Router B> ip bgp peer 129.213.1.1 remote-as 100
Router B> ip bgp peer 175.114.1.2 create
Router B> ip bgp peer 175.114.1.2 remote-as 200
Router B> ip bgp status enable
```

3. On **Router C** assign an AS number of 200, create an external peer of Router D in AS 300, create an internal peer of Router B in AS 200, and enable BGP protocol.

```
Router C> ip bgp as 200
Router C> ip bgp peer 110.210.1.1 create
Router C> ip bgp peer 110.210.1.1 remote-as 100
Router C> ip bgp peer 175.114.1.1 create
Router C> ip bgp peer 175.114.1.1 remote-as 200
Router C> ip bgp status enable
```

4. On **Router D** assign an AS number of 300, create an external peer of Router C in AS200, and enable BGP protocol.

```
Router D> ip bgp as 300
Router D> ip bgp peer 110.210.1.2 create
Router D> ip bgp peer 110.210.1.2 remote-as 200
Router D> ip bgp status enable
```

BGP Policies

BGP policies allow you to define criteria for accepting or rejection routes based on matching aspath lists, prefix lists, or community lists. You can also define route maps that allow you to filter routes based on the same aspath, prefix, or community lists as well as regular expression, prefix, or community primitives. Route maps also allow you to define set values and actions for local preference, community string, MED, origin, and weight parameters when a match is made.

Policy Matching Rules

The OmniCore implementation of BGP follows these rules when accepting or rejection routes based on policy configuration.

1. If no policy is configured, accept the route.
2. If there is no neighbor aspath list configured, goto step 3.

Apply neighbor aspath list match.

If route is rejected, drop it

If route is accepted

 If there is route-map configured go to route-map

 If there is no route-map configured accept route.

If route falls-through-the-cracks-because-there-was-no-catchAll, fall through to step 3.

3. If there is no neighbor prefix list configured, goto step 4

Apply neighbor prefix list match

If route is rejected, drop it

If route is accepted

 If there is route-map configured go to route-map

 If there is no route-map configured accept route.

If route falls-through-the-cracks-because-there-was-no-catchAll fall through to step 4.

4. If there is no neighbor community-list configured, go to route-map

Apply neighbor community list match

If route is rejected, drop it

If route is accepted

If there is route-map configured go to route-map

If there is no route-map configured accept route.

If route falls-through-the-cracks-because-there-was-no-catchAll :

route-map :

next-route-map-instance :

- a.** If no route-map aspath-list is configured, go to step (b)
 apply route-map aspath-list
 If we accept the route, go to step (b)
 If we deny the route, drop it.
 If no match, go to next-route-map-instance.
- b.** If no route-map prefix-list is configured, go to step (c)
 apply route-map prefix-list
 If we accept the route, go to step (c)
 If we deny the route, drop it.
 If no match, go to next-route-map-instance.
- c.** If no route-map community list is configured, go to step (d)
 apply route-map community list
 If we accept the route, go to step (d)
 If we deny the route, drop it.
 If no match, go to next-route-map-instance.
- d.** If regular expression matching primitive is NOT configured, go to step (e)
 Do the match.
 If match, go to step (e)
 If no match, go to next-route-map-instance
- e.** If prefix match primitive is NOT configured, go to step (f)
 Do the match.
 If match, go to step (f)
 If no match, go to next-route-map-instance
- f.** If community matching primitive is NOT configured, go to next-route-map-instance.
 Do the match.
 If match, go to route-map-hit
 If no match, go to next-route-map-instance.

nothing-matches :

reject route

route-map-hit :

if route-map says permit, take set actions and accept the route

if route-map says reject, drop the route

Displaying BGP Policies

You can view current BGP policies by issuing the *ip bgp policy aspath-list show*, *ip bgp policy prefix-list show*, or *ip bgp policy community-list show* command.

```
OmniCore> ip bgp policy aspath-list show
                                     BGP AS Path Lists
-----
AS Path Match List Name      :asListA
AS Path Regular Expression   :400 .* 600$
Priority                      :1
Allow a Matching Route      :permit
Admin Status                 :enable
-----
AS Path Match List Name      :asListA
AS Path Regular Expression   :(500 | 400) ? 300$
Priority                      :1
Allow a Matching Route      :permit
Admin Status                 :enable
-----
Number of Entries Displayed: 2
```

You can view current BGP policy route-maps by issuing the *ip bgp policy route-map show* command.

```
OmniCore> ip bgp policy route-map show
BGP Route Maps
-----
Route Map Name                :BLOCK_ALL
Route Map Sequence Number     :1
Match Prefix                  :255.255.255.255
Match Mask                    :255.255.255.255
Allow Matching Route to Pass :permit
Weight                        :0
Admin Status                  :enable
-----
Route Map Name                :MAP_10
Route Map Sequence Number     :196
Match Prefix                  :196.0.0.0
Match Mask                    :255.0.0.0
Allow Matching Route to Pass :permit
Weight                        :0
Admin Status                  :enable
-----
Route Map Name                :MAP_10
Route Map Sequence Number     :203
Match Prefix                  :203.0.0.0
Match Mask                    :255.0.0.0
Allow Matching Route to Pass :permit
Weight                        :0
Admin Status                  :enable
-----
Number of Entries Displayed: 3
```

BGP Policy Commands

The BGP policy commands are summarized below.

BGP Policy Commands

Command	Default	Description
ip bgp policy aspath-list	no default	Creates, deletes, enables, or disables an AS path list.
ip bgp policy aspath-list action	permit	Sets the action to be taken for an AS path list when a match is found
ip bgp policy aspath-list priority	1	Sets the priority of processing among various regular expressions in an AS path list
ip bgp policy community-list	no default	Creates, deletes, enables, or disables a community list.
ip bgp policy community-list action	permit	Sets the action to be taken for a community list when a match is found.
ip bgp policy community-list match-type	exact	Sets the type of matching to be performed with a community list string.
ip bgp policy community-list priority	1	Sets the priority of processing among multiple items in a community list filter.
ip bgp policy prefix-list	no default	Creates, deletes, enables, or disables a prefix match list.
ip bgp policy prefix-list action	permit	Sets the action to be taken for a prefix list when a match is found
ip bgp policy prefix-list ge	zero (no limit)	Sets the lower limit on the length of the prefix to be matched.
ip bgp policy prefix-list le	zero (no limit)	Sets the upper limit on the length of the prefix to be matched.
ip bgp policy route-map	no default	Creates, deletes, enables, or disables a policy route map.
ip bgp policy route-map action	permit	Sets the action to be taken for a route when a match is found.
ip bgp policy route-map aspath-list	none	Assigns an AS path matching list to the route-map.
ip bgp policy route-map asprepend	none	Sets the aspath prepend action to be taken when a match is found.
ip bgp policy route-map community	none	Sets the action on the community attribute to be taken when a match is found.
ip bgp policy route-map community-list	none	Assigns a community matching list to the route-map.
ip bgp policy route-map community-mode	add	Sets the action to be taken for a community string when a match is found.

BGP Policy Commands (Continued)

ip bgp policy route-map lpref	zero	Sets the local preference value for the route map.
ip bgp policy route-map lpref-mode	none	Sets the action to be taken when setting the local preference attribute for a matching route.
ip bgp policy route-map match-community	none	Specifies a matching community primitive for the route map.
ip bgp policy route-map match-mask	0.0.0.0	Specifies a matching mask primitive in the route map.
ip bgp policy route-map match-prefix	none	Specifies a matching prefix primitive in the route map.
ip bgp policy route-map match-regexp	none	Specifies an AS path matching regular expression primitive in the route map.
ip bgp policy route-map med	zero	Sets the MED (Multi-Exit Discriminator) value for the route map.
ip bgp policy route-map med-mode	none	Sets the action to be taken when setting the MED (Multi-Exit Discriminator) attribute for a matching route.
ip bgp policy route-map origin	none	Sets the action to be taken on the origin attribute when a match is found.
ip bgp policy route-map prefix-list	none	Assigns a prefix matching list to the route-map.
ip bgp policy route-map weight	zero	Sets a BGP weight value to be assigned inbound routes when a match is found.

BGP Redistribution Filters

BGP redistribution is a common method of advertising internal route on the internet. The redistribution filter commands allow you to redistribute your IGP routes into BGP by specifying a source protocol and a destination IP address. Special care must be taken to avoid advertising routes that have been learned through BGP and, therefore, do not need to be sent out to the internet.

You can view current BGP redistribution filters by issuing the *ip bgp redist-filter show* command.

```
OmniCore> ip bgp redist-filter show
Proto      DestinationIP      Mask                Subnets  Effect  RowStat
-----
static     100.10.0.0         255.255.0.0        disable   permit  enable
-----
ospf       100.10.0.0         255.255.0.0        disable   permit  enable
-----
Number of Entries Displayed: 2
```

The BGP redistribution filter commands are summarized below.

BGP Redistribution Filter Commands

Command	Default	Description
ip bgp redist-filter	no default	Creates or deletes a redistribution filter.
ip bgp redist-filter community	“none”	Sets the community string attribute for the redistribution filter.
ip bgp redist-filter effect	permit	Specifies the redistribution filter action for route importation.
ip bgp redist-filter lpref	zero	Sets the local preference attribute value for the redistribution filter.
ip bgp redist-filter metric	zero	Sets the MED attribute value for the redistribution filter.
ip bgp redist-filter subnets	disable	Enables or disables subnet redistribution.