



BGP Address-Family (IPv4/IPv6) Configuration Mode Commands

The Border Gateway Protocol (BGP) Address-Family (IPv4/IPv6) Configuration Mode is used to configure the IPv4 and IPv6 address family information.

Command Modes

Exec > Global Configuration > Context Configuration > BGP Configuration > BGP Address-Family Configuration

configure > **context** *context_name* > **router bgp** *as_number* > **address-family** *address_family_type*

Entering the above command sequence results in the following prompt:

```
[context_name]host_name(config-bgp-af-v6)#
```



Important

The commands or keywords/variables that are available are dependent on platform type, product version, and installed license(s).

- [end](#), on page 1
- [exit](#), on page 2
- [maximum-paths](#), on page 2
- [neighbor](#), on page 3
- [network](#), on page 7
- [redistribute](#), on page 8
- [timers bgp](#), on page 9

end

Exits the current configuration mode and returns to the Exec mode.

Product

All

Privilege

Security Administrator, Administrator

Syntax Description

end

Usage Guidelines Use this command to return to the Exec mode.

exit

Exits the current mode and returns to the parent configuration mode.

Product All

Privilege Security Administrator, Administrator

Syntax Description `exit`

Usage Guidelines Use this command to return to the parent configuration mode.

maximum-paths

Controls the maximum number of parallel external BGP (eBGP) or internal BGP (iBGP) routes that can be installed in a routing table.

Product All

Privilege Security Administrator, Administrator

Command Modes Exec > Global Configuration > Context Configuration > BGP Configuration > BGP Address-Family Configuration

configure > context *context_name* > **router bgp** *as_number* > **address-family** *address_family_type*

Entering the above command sequence results in the following prompt:

```
[context_name]host_name(config-bgp-af-v6) #
```

Syntax Description [**no**] **maximum-paths** { **ebgp** *num_paths* | **ibgp** *num_paths* }
no maximum-paths { **ebgp** | **ibgp** }

no

Disables maximum paths for the specified route type command.

maximum-paths ebgp num_paths

Specifies the maximum number of parallel External Border Gateway Protocol routes as an integer from 1 through 10.

maximum-paths ibgp num_paths

Specifies the maximum number of parallel Internal Border Gateway Protocol routes as an integer from 1 through 10.

**Important**

If configured under the router-bgp-mode, multipath is enabled only for the prefixes learnt in the default-vrf. There is no support for vpnv4 prefixes even though multipath is turned on for the default-vrf.

If configured under the address-family-vrf-mode, multipath is enabled only for prefixes learnt in the vrf.

Usage Guidelines

Use this command to forward packets over multiple paths. User can control the maximum number of parallel eBGP routes that can be installed in a routing table. Enabling multipath does not affect the best path selection in BGP. If multipath is enabled, all the paths with the same weight, local-preference, as-path length, origin, and multi-exit discriminator (MED) as the best path are added to the routing table.

Example

The following command disables forward of packets over multiple paths:

```
no maximum-paths ebgp
```

neighbor

Configures the IPv4/IPv6 Address Family for BGP routers that interconnect to non-broadcast networks.

Product

All

Privilege

Security Administrator, Administrator

Command Modes

Exec > Global Configuration > Context Configuration > BGP Configuration > BGP Address-Family Configuration

```
configure > context context_name > router bgp as_number > address-family address_family_type
```

Entering the above command sequence results in the following prompt:

```
[context_name]host_name(config-bgp-af-v6)#
```

Syntax Description

```
[ no ] neighbor ip_address { activate | advertisement-interval adv_time |
capability graceful-restart | default-originate [ route-map map_name ] |
distribute-list dist_list { in | out } | ebgp-multihop [ max-hop number ]
encrypted password encryp_password fall-over bfd multihop | filter-list
filt_list { in | out } | max-prefix max_num [ threshold thresh_percent ] [
warning-only ] next-hop-self | password password | remote-as AS_num |
remove-private-AS | restart-time rest_time | route-map map_name { in | out }
| send-community { both | extended | standard } | shutdown |
srp-activated-soft-clear | timers { connect-interval connect_interval [
keepalive-interval keepalive_interval holdtime-interval holdtime_interval [
min-peer-holdtime-interval min_peer_hold_interval ] ] | keepalive-interval
keepalive_interval holdtime-interval holdtime_interval { connect-interval
connect_interval | min-peer-holdtime-interval min_peer_hold_interval [
connect-interval connect_interval ] } } | update-source ip_address | weight
value }
```

no

Delete the specified parameter from the router configuration.

neighbor *ip_address*

Specifies the IP address of a BGP neighbor. *ip_address* must be in IPv4 dotted-decimal or IPv6 colon-separated-hexadecimal notation.

activate

Enables the exchange of routes with this neighbor.

advertisement-interval *adv_time*

Specifies the minimum interval (in seconds) between sending BGP routing updates. *adv_time* must be an integer from 0 through 600. Default: 30

capability graceful-restart

Configures graceful re-start attributes.

default-originate [route-map *map_name*]

Specifies the default originate routes to this neighbor

route-map *map_name*: Specifies the route-map that contains the criteria to originate default routes. *map_name* must be the name of an existing route-map in the current context.

distribute-list *dist_list* { in | out }

Filters updates to and from this neighbor based on a route access list. Default: No filtering is performed. *dist_list* is the name or number of an existing route-access-list.

in: Indicates that incoming advertised routes should be filtered.

out: Indicates that outgoing advertised routes should be filtered.

ebgp-multihop [max-hop *number*]

Allows eBGP neighbors that are not on directly connected networks.

max-hop *number*: Specifies the maximum number of hops allowed to reach a neighbor as an integer from 1 through 255. Default hop count: 255

encrypted password *encryp_password*

Specifies the encrypted password that is used only inside configuration files. This is an alphanumeric string of 1 through 24 characters.

fall-over bfd multihop

Supports Bidirectional Forwarding Detection (BFD) multihop for fallover.

filter-list *filt_list* { in | out }

Establishes BGP filters based on an autonomous system (AS) path access list. *filt_list* is name of an existing AS path access list.

in: Indicates that incoming advertised routes will be filtered.

out: Indicates that outgoing advertised routes will be filtered.

max-prefix *max_num* [threshold *thresh_percent*] [warning-only]

Specifies the maximum number of prefixes accepted from this peer. When the maximum is exceeded the neighbor connection is reset. *max_num* is an integer from 1 through 4294967295. Default: No maximum prefix limit.

threshold *thresh_percent*: Specifies a percentage value of when the BGP table is full. When this value is reached peer warnings are sent to the neighbor. *thresh_percent* must be an integer from 1 through 100.

warning-only: Specifies that only a warning message is sent when the limit is exceeded. The neighbor connection is not reset

next-hop-self *ip_address*

Disables the next hop calculation for this neighbor.

password *password*

Sets a *password* expressed as an alphanumeric string of 1 through 24 characters.

remote-as *AS_num*

Specify the AS number of the BGP neighbor as an integer from 1 through 4294967295.

remove-private-AS

Removes the private AS number from outbound updates. Default: Do not remove the private AS number.

restart-time *rest_time*

Specifies the maximum time (in seconds) required to for neighbor to restart as an integer from 1 through 3600.

route-map *map_name* { in | out }

Applies a route map to the neighbor. *map_name* must be the name of an existing route-map in the current context.

in: Indicates that the route map applies to incoming advertisements.

out: Indicates that the route map applies to outgoing advertisements.

send-community { both | extended | standard }

Sends the community attributes to a peer router (neighbor).

both: Sends standard and extended community attributes

extended: Sends extended community attributes.

standard:Sends standard community attributes.

shutdown

Administratively shuts down this neighbor. This disables exchanging routes or configuring parameters for this neighbor.

srp-activated-soft-clear

Enables BGP updates when SRP-enabled resources are modified.

timers { [**connect-interval** *connect_interval*] [[**keepalive-interval** *keepalive_interval* / **holdtime-interval** *holdtime_interval*]] }

Sets BGP timers for the specified neighbor.

connect-interval *connect_interval*: Specifies the connect timer (in seconds) as an integer from 0 through 65535. The default is 60 seconds.

keepalive-interval *keepalive_interval*: Specifies the frequency (in seconds) at which the current BGP router sends keepalive messages to its neighbor. *keep_time* must be an integer from 0 through 65535. The default is 30 seconds.

holdtime-interval *holdtime_interval*: Specifies the interval (in seconds) the router waits for a keepalive message before declaring a neighbor dead. *hold_time* must be an integer from 0 through 65535. The default is 90 seconds.

min-peer-holdtime-interval *min_peer_hold_interval*: Specifies the minimum acceptable hold time (in seconds) from peer for a keepalive message before declaring a neighbor dead. *min_peer_hold_interval* must be an integer from 0 through 65535. The default is 90 seconds.

update-source *ip_address*

Binds the specified IP address to the BGP socket that is used to communicate to the peer. *ip_address* is an IPv4 address in dotted-decimal notation.

In most cases you should set the update-source address to the address of the loopback interface in the current context. By doing this, the TCP connection does not go down until there is no route for the loopback address in the peering router.

weight *value*

Sets the default weight for routes from this neighbor as an integer from 0 through 65535. Default: 0

Usage Guidelines

Use this command to set parameters for communication with a specified neighbor. The chassis supports a maximum of 64 peers per context.



Important

A remote AS number must be specified for a neighbor before other parameters can be configured.

Example

The following command specifies that the neighbor at the IP address *192.168.100.25* has an AS number of *2000*:

```
neighbor 192.168.100.25 remote-as 2000
```

The following command allows BGP neighbors that are a maximum of 27 hops away:

```
neighbor 192.168.100.25 ebgp-multihop max-hop 27
```

The following command sets the minimum interval between sending routing updates to 3 minutes (180 seconds):

```
neighbor 192.168.100.25 advertisement-interval 180
```

The following command sets the default weight for all routes from the specified neighbor to *100*:

```
neighbor 192.168.100.25 weight 100
```

network

Configures and specifies a network to announce via BGP.

Product

All

Privilege

Security Administrator, Administrator

Command Modes

Exec > Global Configuration > Context Configuration > BGP Configuration > BGP Address-Family Configuration

configure > **context** *context_name* > **router bgp** *as_number* > **address-family** *address_family_type*

Entering the above command sequence results in the following prompt:

```
[context_name]host_name(config-bgp-af-v6)#
```

Syntax Description

[**no**] **network** *ip_address/mask* [**route-map** *map_name*]

no

Delete the specified network from the configuration for the BGP router.

network *ip_address/mask*

Specifies the IP address and netmask bits for the network to announce via BGP. *ip_address* is a network IP address in IPV4 dotted-decimal notation and *mask* is the number of subnet bits, representing a subnet mask in CIDR. These must be entered in the IPV4 dotted-decimal notation/subnet bits format.

route-map *map_name*

Filter routes through the specified route map before announcing the network. *map_name* specifies the name of the route-map to use as an alphanumeric string of 1 through 79 characters.

Usage Guidelines

Use this command to specify a network to announce via BGP.

Example

The following command announces the network *192.168.0.0* with a netmask of *16* via BGP:

```
network 192.168.0.0/16
```

The following command removes the network from the BGP router configuration:

```
no network 192.168.0.0/16
```

redistribute

Redistributes routes into BGP from another protocol as BGP neighbors.

Product

All

Privilege

Security Administrator, Administrator

Command Modes

Exec > Global Configuration > Context Configuration > BGP Configuration > BGP Address-Family Configuration

```
configure > context context_name > router bgp as_number > address-family address_family_type
```

Entering the above command sequence results in the following prompt:

```
[context_name]host_name(config-bgp-af-v6) #
```

Syntax Description

```
[ no ] redistribute { connected | ospf | rip | static } [ route-map map_name ]
```

no

Remove the specified redistribution parameters from the BGP router configuration.

redistribute connected

Specifies that connected routes will be redistributed.

redistribute ospf

Specifies that Open Shortest Path First (OSPF) routes will be redistributed.

redistribute rip

Specifies that Routing Information Protocol (RIP) routes will be redistributed. (RIP is not supported at this time.)

redistribute static

Specifies that static routes will be redistributed.

route-map *map_name*

Filters routes through the specified route map before redistribution. *map_name* specifies the name of the route-map to use as an alphanumeric string of 1 through 79 characters.

Usage Guidelines

Use this command to specify what routes this BGP router should redistribute into BGP.

Example

The following command redistributes OSPF routes after filtering them through the route map named *Map1*:

```
redistribute ospf route-map Map1
```

The following command removes the redistribution of OSPF routes from the router's configuration:

```
no redistribute ospf route-map Map1
```

timers bgp

Enables or disables an aggressive minimum BGP route advertisement interval (MinRtAdvInterval) for ICSR configurations.

Product

All products that support ICSR

Privilege

Security Administrator, Administrator

Command Modes

Exec > Global Configuration > Context Configuration > BGP Configuration > BGP Address-Family Configuration

```
configure > context context_name > router bgp as_number > address-family address_family_type
```

Entering the above command sequence results in the following prompt:

```
[context_name]host_name(config-bgp-af-v6) #
```

Syntax Description

```
[ no ] timers bgp icshr-aggr-advertisement-interval seconds
```

no

Disables this aggressive ICSR BGP advertisement interval.

seconds

Sets the number of seconds as an integer from 0 to 30. Default: 0.

Usage Guidelines

Use this command to configure an aggressive ICSR BGP advertisement interval (MinRtAdvInterval). The default value is 0. If set as 0, the aggressive advertisement interval is disabled.

The MinRtAdvInterval can be uniquely set for each address family.

After ICSR switchover, BGP will set the advertisement-interval for each AFI/SAFI (Address Family Identifier/Subsequent Address Family Identifier) supported by the peer to the configured value. BGP updates will be advertised to the peer based on this interval.

Example

The following command sets the MinRtAdvInterval for this address family to 1 second:

```
timers bgp icsr-aggr-advertisement-interval 1
```