N Commands

This chapter describes the Cisco NX-OS Border Gateway Protocol (BGP) commands that begin with N.

neighbor

To configure a Border Gateway Protocol (BGP) neighbor (router or VRF) and enter the neighbor configuration mode, use the **neighbor** command. To remove an entry, use the **no** form of this command.

neighbor {*ip-addr* | *ip-prefix/length*} [**remote-as** {*as-num*[.*as-num*] | **route-map** *name*}

no neighbor {*ip-addr* | *ip-prefix/length*} [**remote-as** {*as-num*[.*as-num*] | **route-map** *name*}]

Syntax Description

ip-addr	IP address of the neighbor in this format: A.B.C.D.
ip-prefix/length	IP prefix and the length of the IP prefix. The format is x.x.x.x/*length*. The *length* range is from 1 to 32.
remote-as	(Optional) Specifies the autonomous system (AS) number of the neighbor.
as-num	Number of an AS that identifies the router to other BGP routers and tags the routing information passed along. The range is from 1 to 65535.
.*as-num*	(Optional) Number of an AS that identifies the router to other BGP routers and tags the routing information passed along. The range is from 1 to 65535.
route-map *name*	(Optional) Specifies a route map that matches the BGP peer AS number against a list of AS numbers or a regular expression. The name can be any case-sensitive, alphanumeric string up to 63 characters.

Command Default

None

Command Modes

Neighbor address-family configuration mode

Router BGP configuration mode

Command History

Release	Modification
6.0(2)N1(1)	This command was introduced.

Usage Guidelines

From the BGP neighbor configuration mode, you can perform the following actions:

- **address-family**—Configures an address-family (router, neighbor, VRF). See the **address-family (BGP)** command for information.
- **description** *description*—Describes the neighbor. You can enter up to 80 characters including spaces.
- **disable-connected-check**—Disables the connection verification for the directly connected peer. Use the **disable-connected-check** command to disable a check for an exterior Border Gateway Protocol (eBGP) peer that is directly connected to the local router. BGP triggers a connection check

automatically for all eBGP peers that are known to be a single hop away, unless you disable this check with the **disable-connected-check** command. BGP does not bring up sessions if the check fails. BGP considers an EBGP peer as a single hop away if the eBGP peer does does not have the **ebgp-multihop** command configured (that is, the time-to-live (TTL) value is one).

This command is ignored if the **route-map** keyword is used in the **neighbor** command.

- **dont-capability-negotiate**—Turns off the negotiate capability with this neighbor.
- **dynamic-capability**—Enables the dynamic capability.
- **ebgp-multihop**—Accepts and attempts BGP connections to external peers that reside on networks that are not directly connected. This command is ignored if the **route-map** keyword is used in the **neighbor** command.

Note You should enter this command under the guidance of Cisco technical support staff only.

- **exit**—Exits from the current command mode.
- **inherit peer-session** *session-name*—Configures a peer to inherit the configuration from another peer-session template. To remove an inherit statement from a peer-session template, use the **no** form of this command.
- **no**—Negates a command or sets its defaults.
- **transport connection-mode passive**—Allows a passive connection setup only. To remove the restriction, use the **no** form of this command.
- **remove-private-as**—Removes the private AS number from the outbound updates.
- **shutdown**—Administratively shuts down this neighbor.
- **timers** *keepalive-time*—Configures keepalive and hold timers in seconds. The range is from 0 to 3600. The default is 60.
- **update-source** {**ethernet** *mod/port* | **loopback** *virtual-interface* | **port-channel** *number*[**.***sub-interface*]}—Specifies the source of the BGP session and updates. The range for *virtual-interface* is from 0 to 1023. The range for *number* is from 0 to 4096. The range for *sub-interface* is from 1 to 4093.

  The Cisco NX-OS software allows BGP sessions to use any operational interface for TCP connections when you enter the **update-source** command in neighbor configuration mode. To restore the interface assignment to the closest interface, which is called the best local address, use the **no** form of this command.

This command requires the LAN Enterprise Services license.

Examples

This example shows how to configure a single-hop eBGP peering session between two BGP peers that are reachable on the same network segment through a local loopback interfaces on each router:

BGP Peer 1

```
switch(config)# interface loopback 1
switch(config-if)# ip address 10.0.0.100 255.255.255
switch(config-if)# exit
switch(config)# router bgp 64497
switch(config-router)# neighbor 192.168.0.200 remote-as 64496
switch(config-router-neighbor)# update-source loopback 2
switch(config-router-neighbor)# disable-connected-check
switch(config-router-neighbor)#
```

BGP Peer 2

```
switch(config)# interface loopback 2
switch(config-if)# ip address 192.168.0.200 255.255.255
switch(config-if)# exit
switch(config)# router bgp 64496
switch(config-router)# neighbor 10.0.0.100 remote-as 64497
switch(config-router-neighbor)# update-source loopback 1
switch(config-router-neighbor)# disable-connected-check
switch(config-router-neighbor)#
```

This example shows how to source BGP TCP connections for the specified neighbor with the IP address of the loopback interface rather than the best local address:

```
switch(config)# router bgp 64496
switch(config-router)# neighbor 172.16.0.0 remote-as 64496
switch(config-router-neighbor)# update-source Loopback0
switch(config-router-neighbor)#
```

Related Commands

Command	Description
feature bgp	Enables BGP on the router.
route-map	Creates a route map.

network

To configure an IP prefix to advertise, use the **network** command. To remove the IP prefix to advertise, use the **no** form of this command.

network *ip-addr* | *ip-prefix/length* **mask** *mask-num* [**route-map** *name*]

no network *ip-network* | *ip-prefix/length* **mask** *mask-num* [**route-map** *name*]

Syntax Description

ip-addr	IP network address to advertise; use the following format: A.B.C.D.
ip-prefix/length	IP prefix and the length of the IP prefix. Use the following format: A.B.C.D/length.
mask *mask-num*	Configures the mask of the IP prefix to advertise in dotted 4-octet format.
route-map *name*	(Optional) Specifies the name of the route map to modify attributes.

Command Default

None

Command Modes

Neighbor address-family configuration mode

Router BGP configuration mode

Command History

Release	Modification
6.0(2)N1(1)	This command was introduced.

Usage Guidelines

The IP prefix to advertise is considered as a best path and advertisement to peers only if a route of equal or more specificity is present in the routing table.

Examples

This example shows how to configure an IP prefix to advertise:

```
switch(config-router-af)# network 2.2.2.2 mask 3.3.3.3 route-map test
switch(config-router-af)#
```

Related Commands

Command	Description
show ip prefix-list	Displays information about IP prefix lists.

nexthop route-map

To specify that Border Gateway Protocol (BGP) routes are resolved using only the next hops that have routes that match specific characteristics, use the **nexthop route-map** command. To remove the route map, use the **no** form of this command.

nexthop route-map *name*

no nexthop route-map *name*

Syntax Description

name	Route map name. The name can be any alphanumeric string up to 63 characters.

Command Default

None

Command Modes

Adress family configuration mode

Command History

Release	Modification
6.0(2)N1(1)	This command was introduced.

Usage Guidelines

Use the **nexthop route-map** command to configure route policy filtering for next hops.

BGP next-hop filtering allows you to specify that when a next-hop address is checked with the Routing Information Base (RIB), the underlying route for that next-hop address is passed through the route map. If the route map rejects the route, the next-hop address is treated as unreachable.

BGP marks all next hops that are rejected by the route policy as invalid and does not calculate the best path for the routes that use the invalid next-hop address.

This command requires an LAN Enterprise Services license.

Examples

This example shows how to configure a route map to filter the next-hop address:

```
switch# configure terminal
switch(config)#route-map CHECK-BGP25 deny 10
switch(config-route-map)# match ip address prefix-list FILTER25
switch(config-route-map)# match source-protocol ospf-o1
switch(config-route-map)# exit
switch(config)#ip prefix-list FILTER25 seq 5 permit 0.0.0.0/0 le 25
switch(config)# router bgp 1.0
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# nexthop route-map CHECK-BGP25
switch(config-router-af)#
```

Related Commands

Command	Description
feature bgp	Enables BGP.
nexthop trigger-delay	Configures the delay timers for BGP next-hop address tracking.
route-map	Defines the conditions for redistributing routes from one routing protocol into another.

next-hop-self

To set the IP address of the router as the next hop address, use the **next-hop-self** command. To revert to the default configuration, use the **no** form of this command.

next-hop-self

no next-hop-self

Syntax Description

This command has no arguments or keywords.

Command Default

None

Command Modes

Address family configuration mode

Command History

Release	Modification
6.0(2)N1(1)	This command was introduced.

Usage Guidelines

This command requires a LAN Enterprise Services license.

Examples

This example shows how to configure the IP address of a router as the next-hop address:

```
switch# configure terminal
switch(config)# router bgp 102
switch(config-router)# neighbor 192.168.1.3 remote-as 64497
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# next-hop-self
switch(config-router-neighbor-af)#
```

Related Commands

Command	Description
address-family (BGP neighbor)	Enters the BGP neighbor address-family configuration mode.
feature bgp	Enables BGP.
show ip bgp	Displays BGP configuration information.

nexthop trigger-delay

To specify a Border Gateway Protocol (BGP) delay for triggering next-hop calculations, use the **nexthop trigger-delay** command. To set the trigger delay to the default value, use the **no** form of this command.

nexthop trigger-delay {**critical** *delay* | **non-critical** *delay*}

no nexthop trigger-delay {**critical** *delay* | **non-critical** *delay*}

Syntax Description

critical *delay*	Specifies the critical next-hop trigger delay, in milliseconds. The range is from 0 to 4294967295. The default is 3000.
non-critical *delay*	Specifies the noncritical next-hop trigger delay, in milliseconds. The range is from 0 to 4294967295. The default is 10000.

Command Default

Critical delay: 3000 milliseconds.
Noncritical delay: 10000 milliseconds.

Command Modes

Address family configuration mode

Command History

Release	Modification
6.0(2)N1(1)	This command was introduced.

Usage Guidelines

Use the **nexthop trigger-delay** command to modify when BGP processes next-hop address tracking events.

The **non-critical** *delay* value must always be set to a value that is at least equal or greater to the **critical** *delay* value.

The delay should be slightly higher than the time it takes for the Interior Gateway Protocol (IGP) to settle into a steady state after some event (IGP convergence time).

This command requires a LAN Enterprise Services license.

Examples

This example shows how to modify the next-hop address tracking delay:

```
switch# configure terminal
switch(config)# router bgp 1.0
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# nexthop trigger-delay critical 5000 non-critical 20000
```

Related Commands

Command	Description
feature bgp	Enables BGP.
nexthop route-map	Configures a route map for BGP next-hop address tracking.