



BGP Commands

This chapter describes the commands used to configure and monitor Border Gateway Protocol (BGP) IP Version 4 (IPv4), IP Version 6 (IPv6),) routing sessions.

For detailed information about BGP concepts, configuration tasks, and examples, see the *Implementing BGP* chapter in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.



Note Running the **show bgp** command immediately after configuring a large and complex route policy may result in timeout of the system database shown through an error message (`SYSDB-SYSDB-6-TIMEOUT_EDM`). It is recommended that the show command be run after the new route policy takes effect.

- [address-family \(BGP\)](#), on page 6
- [advertisement-interval](#), on page 8
- [af-group](#), on page 10
- [aggregate-address](#), on page 12
- [aigp](#), on page 14
- [aigp send-cost-community](#), on page 15
- [allocate-label](#), on page 17
- [allow vpn default-originate](#), on page 18
- [allowas-in](#), on page 19
- [as-format](#), on page 20
- [as-path-loopcheck out disable](#), on page 21
- [attribute-filter group](#), on page 22
- [attribute large-community](#), on page 23
- [bfd \(BGP\)](#), on page 24
- [bgp as-path-loopcheck](#), on page 26
- [bgp attribute-download](#), on page 27
- [bgp auto-policy-soft-reset disable](#), on page 29
- [bgp bestpath as-path ignore](#), on page 30
- [bgp bestpath compare-routerid](#), on page 32
- [bgp bestpath cost-community ignore](#), on page 33
- [bgp bestpath med always](#), on page 34
- [bgp bestpath med confed](#), on page 36
- [bgp bestpath med missing-as-worst](#), on page 38

- [bgp bestpath origin-as allow invalid](#), on page 39
- [bgp bestpath origin-as use validity](#), on page 41
- [bgp bestpath aigp ignore](#), on page 42
- [bgp bestpath as-path multipath-relax](#) , on page 43
- [bgp client-to-client reflection disable](#), on page 44
- [bgp cluster-id](#), on page 46
- [bgp confederation identifier](#), on page 48
- [bgp confederation peers](#), on page 50
- [bgp dampening](#), on page 52
- [bgp default local-preference](#), on page 54
- [bgp enforce-first-as disable](#), on page 55
- [bgp fast-external-fallover disable](#), on page 57
- [bgp graceful-restart](#), on page 58
- [bgp graceful-restart graceful-reset](#), on page 60
- [bgp graceful-restart purge-time](#), on page 61
- [bgp graceful-restart restart-time](#), on page 62
- [bgp graceful-restart stalepath-time](#), on page 63
- [bgp label-delay](#), on page 65
- [bgp log neighbor changes disable](#), on page 66
- [bgp maximum neighbor](#), on page 68
- [bgp nexthop resolution allow-default](#), on page 69
- [bgp redistribute-internal](#), on page 70
- [bgp router-id](#), on page 72
- [bgp scan-time](#), on page 73
- [bgp update-delay](#), on page 74
- [capability additional-paths receive](#), on page 75
- [capability additional-paths send](#), on page 76
- [capability orf prefix](#), on page 77
- [capability suppress](#) , on page 79
- [clear bgp](#), on page 82
- [cef consistency-hashing auto-recovery](#), on page 84
- [clear bgp dampening](#), on page 85
- [clear bgp external](#), on page 87
- [clear bgp flap-statistics](#), on page 88
- [clear bgp nexthop performance-statistics](#), on page 90
- [clear bgp nexthop registration](#), on page 92
- [clear bgp peer-drops](#), on page 94
- [clear bgp performance-statistics](#), on page 95
- [clear bgp rpki](#) , on page 96
- [clear bgp self-originated](#), on page 98
- [clear bgp shutdown](#), on page 100
- [clear bgp soft](#), on page 101
- [cluster-id](#), on page 104
- [cluster-id allow-equal](#) , on page 106
- [default-information originate \(BGP\)](#), on page 107
- [default-metric \(BGP\)](#), on page 108

- [default-originate](#), on page 109
- [description \(BGP\)](#), on page 111
- [distance bgp](#), on page 112
- [dscp \(BGP\)](#), on page 114
- [ebgp-multihop](#), on page 116
- [enforce-first-as](#), on page 118
- [enforce-first-as-disable](#), on page 120
- [ibgp policy out enforce-modifications](#), on page 122
- [ignore-connected-check](#), on page 123
- [keychain](#), on page 124
- [keychain-disable](#), on page 126
- [keychain inheritance-disable](#), on page 128
- [local-as](#), on page 130
- [maximum-paths \(BGP\)](#), on page 132
- [maximum-prefix \(BGP\)](#), on page 134
- [mpls activate \(BGP\)](#), on page 137
- [multipath](#), on page 138
- [neighbor \(BGP\)](#), on page 139
- [neighbor-group](#), on page 141
- [network \(BGP\)](#), on page 143
- [network backdoor](#), on page 145
- [next-hop-self](#), on page 147
- [next-hop-unchanged](#), on page 149
- [nexthop route-policy](#), on page 151
- [nexthop trigger-delay](#), on page 153
- [nsr \(BGP\)](#), on page 155
- [nsr disable \(BGP\)](#), on page 157
- [orf](#), on page 159
- [password \(BGP\)](#), on page 161
- [password-disable](#), on page 163
- [precedence](#), on page 165
- [receive-buffer-size](#), on page 167
- [redistribute \(BGP\)](#), on page 169
- [remote-as \(BGP\)](#), on page 173
- [remove-private-as](#), on page 176
- [route-policy \(BGP\)](#), on page 178
- [route-reflector-client](#), on page 180
- [router bgp](#), on page 182
- [send-buffer-size](#), on page 183
- [send-community-ebgp](#), on page 185
- [send-extended-community-ebgp](#), on page 187
- [session-group](#), on page 189
- [session-open-mode](#), on page 191
- [show bgp](#), on page 193
- [show bgp large-community](#), on page 203
- [send-community-ebgp](#), on page 205

- [show bgp update out](#), on page 207
- [show bgp update in error process](#), on page 208
- [show bgp update out filter-group](#), on page 209
- [show bgp update out process](#), on page 210
- [show bgp update out sub-group](#), on page 212
- [show bgp update out update-group](#), on page 213
- [show bgp advertised](#), on page 214
- [show bgp af-group](#), on page 220
- [show bgp attribute-key](#), on page 223
- [show bgp cidr-only](#), on page 226
- [show bgp community](#), on page 230
- [show bgp convergence](#), on page 235
- [show bgp dampened-paths](#), on page 238
- [show bgp flap-statistics](#), on page 242
- [show bgp inconsistent-as](#), on page 247
- [show bgp labels](#), on page 251
- [show bgp neighbor-group](#), on page 252
- [show bgp neighbors](#), on page 256
- [show bgp neighbors nsr](#), on page 273
- [show bgp nexthops](#), on page 275
- [show bgp nsr](#), on page 277
- [show bgp paths](#), on page 281
- [show bgp policy](#), on page 283
- [show bgp process](#), on page 290
- [show bgp regexp](#), on page 308
- [show bgp route-policy](#), on page 312
- [show bgp session-group](#), on page 316
- [show bgp sessions](#), on page 319
- [show bgp summary](#), on page 320
- [show bgp summary nsr](#), on page 324
- [show bgp table](#), on page 327
- [show bgp truncated-communities](#), on page 329
- [show bgp update-group](#), on page 333
- [show protocols \(BGP\)](#), on page 336
- [show tcp brief](#), on page 338
- [shutdown \(BGP\)](#), on page 339
- [socket receive-buffer-size](#), on page 341
- [socket send-buffer-size](#), on page 343
- [soft-reconfiguration inbound](#), on page 345
- [speaker-id](#), on page 348
- [table-policy](#), on page 349
- [timers \(BGP\)](#), on page 351
- [timers bgp](#), on page 353
- [ttl-security](#), on page 355
- [update limit](#), on page 358
- [update limit address-family](#), on page 359

- [update limit sub-group, on page 360](#)
- [update in error-handling basic disable, on page 362](#)
- [update in error-handling extended, on page 363](#)
- [update in labeled-unicast equivalent, on page 364](#)
- [update out logging, on page 365](#)
- [update-source, on page 366](#)
- [use, on page 368](#)
- [weight, on page 372](#)

address-family (BGP)

To enter various address family configuration modes while configuring Border Gateway Protocol (BGP), use the **address-family** command in an appropriate configuration mode. To disable support for an address family, use the **no** form of this command.

```
address-family {ipv4 {labeled-unicast | flowspec | unicast} | ipv6 {labeled-unicast | flowspec | unicast}
| |}
no address-family
```

Syntax Description

| | |
|-----------------------------|--|
| ipv4 unicast | Specifies IP Version 4 (IPv4) unicast address prefixes. |
| ipv4 labeled-unicast | Specifies IPv4 labeled-unicast address prefixes. This option is available in IPv4 neighbor configuration mode. |
| ipv6 unicast | Specifies IP Version 6 (IPv6) unicast address prefixes. |
| ipv6 labeled-unicast | Specifies IPv6 labeled-unicast address prefixes. This option is available in IPv6 neighbor configuration mode. |

Command Default

An address family must be explicitly configured in the router configuration mode for the address family to be active in BGP. Similarly, an address family must be configured under the neighbor for the BGP session to be established for that address family. An address family must be configured in router configuration mode before it can be configured under a neighbor.

Command Modes

Router configuration
 Neighbor configuration
 Neighbor group configuration
 Flowspec configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family** command to enter various address family configuration modes while configuring BGP routing sessions. When you enter the **address-family** command from router configuration mode, you enable the address family and enter global address family configuration mode.

The IPv4 unicast address family must be configured in router configuration mode before configuring the IPv4 labeled-unicast address family for a neighbor in neighbor configuration mode. The IPv6 unicast address family must be configured in router configuration mode before configuring the IPv6 labeled-unicast address family for a neighbor in neighbor configuration mode.

When you enter the **address-family** command from neighbor configuration mode, you activate the address family on the neighbor and enter neighbor address family configuration mode. IPv4 neighbor sessions support IPv4 unicast, , and labeled-unicast address families. IPv6 neighbor sessions support IPv6 unicast address families.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to place the router in global address family configuration mode for the IPv4 address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#
```

The following example shows how to place the router in flowspec sub-address family configuration mode for the IPv4 address family:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 flowspec
RP/0/RP0/CPU0:router(config-bgp-af)#
```

advertisement-interval

To set the minimum interval between the sending of Border Gateway Protocol (BGP) routing updates, use the **advertisement-interval** command in an appropriate configuration mode. To remove the **advertisement-interval** command from the configuration file and restore the system to its default interval values, use the **no** form of this command.

advertisement-interval *seconds*

no advertisement-interval [*seconds*]

Syntax Description

seconds Minimum interval between sending BGP routing updates (in seconds). Range is 0 to 600.

Command Default

Default minimum interval:

For internal BGP (iBGP) peers is 0 seconds

For external BGP (eBGP) peers is 30 seconds

For customer edge (CE) peers is 0 seconds

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | No modification. |
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If this command configures a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to set the minimum time between sending BGP routing updates to 10 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 5
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
```



```
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100  
RP/0/RP0/CPU0:router(config-bgp-nbr)# advertisement-interval 10
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

af-group

To create an address family group for Border Gateway Protocol (BGP) neighbors and enter address family group configuration mode, use the **af-group** command in router configuration mode. To remove an address family group, use the **no** form of this command.

af-group *af-group-name* **address-family**
no af-group

| | | |
|---------------------------|-----------------------------|---|
| Syntax Description | <i>af-group-name</i> | Address family group name. |
| | address-family | Enters address family configuration mode. |
| | ipv4 unicast | Specifies IP Version 4 (IPv4) unicast address prefixes. |
| | ipv4 labeled-unicast | Specifies IPv4 labeled unicast address prefixes. |
| | ipv6 unicast | Specifies IP Version 6 (IPv6) unicast address prefixes. |
| | ipv6 labeled-unicast | Specifies IPv6 labeled unicast address prefixes. |

Command Default No BGP address family group is configured.

Command Modes Router configuration

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **af-group** command to group address family-specific neighbor commands within an IPv4 or IPv6 address family. Neighbors that have address family configuration are able to use the address family group. Further, neighbors inherit the configuration parameters of the entire address family group.

You cannot define two address family groups with the same name in different address families.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to create address family group group1 and enter address family group configuration mode for IPv4 unicast. Group1 contains the next-hop-self feature, which is inherited by neighbors that use address family group1.

```
RP/0/RP0/CPU0:router(config)# router bgp 100  
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-self
```

Related Commands

| Command | Description |
|---|---|
| neighbor (BGP), on page 139 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| use, on page 368 | Inherits configuration from a neighbor group, session group, or address family group. |

aggregate-address

To create an aggregate entry in a Border Gateway Protocol (BGP) routing table, use the **aggregate-address** command in an appropriate configuration mode. To remove the **aggregate-address** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

aggregate-address *address/mask-length* [**as-set**] [**as-confed-set**] [**summary-only**] [**route-policy** *route-policy-name*]
no aggregate-address

Syntax Description

| | |
|--|--|
| <i>address</i> | Aggregate address. |
| <i>/mask-length</i> | Aggregate address mask length. |
| as-set | (Optional) Generates autonomous system set path information and community information from contributing paths. |
| as-confed-set | (Optional) Generates autonomous system confederation set path information from contributing paths. |
| summary-only | (Optional) Filters all more-specific routes from updates. |
| route-policy <i>route-policy-name</i> | (Optional) Specifies the name of a route policy used to set the attributes of the aggregate route. |

Command Default

When you do not specify this command, no aggregate entry is created in the BGP routing table.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | No modification. |
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can implement aggregate routing in BGP either by redistributing an aggregate route into BGP using the **network** command or the **aggregate-address** command.

Use the **aggregate-address** command without optional arguments to create an aggregate entry in the BGP routing table if any more-specific BGP routes are available that fall in the specified range. The aggregate route is advertised as coming from your autonomous system and has the atomic aggregate attribute set to show that information might be missing. (By default, the atomic aggregate attribute is set unless you specify the **as-set** keyword.)

Use of the **as-set** keyword creates an aggregate entry using the same rules that the command follows without this keyword. However, the advertised path for this route is an AS_SET, a set of all autonomous systems contained in all paths that are being summarized.

Do not use this form of the **aggregate-address** command when aggregating many paths because this route must be continually withdrawn and updated as autonomous system path reachability information for the summarized routes changes.

Use the **as-confed-set** keyword to create an AS_CONFED_SET in the autonomous system path of the aggregate from any confederation segments in the paths being summarized. This keyword takes effect only if the **as-set** keyword is also specified.

Use of the **summary-only** keyword creates an aggregate entry (for example, 10.0.0.0/8) but suppresses advertisements of more-specific routes to all neighbors. If you want to suppress only advertisements to certain neighbors, use the **route-policy (BGP)** command in neighbor address family configuration mode with caution. If a more-specific route leaks out, all BGP speakers (the local router) prefer that route over the less-specific aggregate you generate (using longest-match routing).

Use the **route-policy** keyword to specify a routing policy for the aggregate entry. The **route-policy** keyword is used to select which more-specific information to base the aggregate entry on and which more-specific information to suppress. You can also use the keyword to modify the attributes of the aggregate entry.

Task ID

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to create an aggregate address. The path advertised for this route is an autonomous system set consisting of all elements contained in all paths that are being summarized.

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# aggregate-address 10.0.0.0/8 as-set
```

Related Commands

| Command | Description |
|---|---|
| network (BGP), on page 143 | Specifies the list of networks for the BGP routing process. |
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |

aigp

To enable sending and receiving of accumulated interior gateway protocol (AiGP) attribute per eBGP neighbor, use the **aigp** command in appropriate configuration mode. To disable this functionality, either use the **disable** keyword or use the **no** form of this command.

aigp [**disable**]

no aigp

Syntax Description

disable Disables sending or receiving AiGP attribute.

Command Default

Send or receive of AiGP attribute is disabled for eBGP neighbors

Command Modes

IPv4 address family configuration

IPv6 address family configuration

Neighbor address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to enable AiGP send and receive capability under neighbor address family (IPv4 unicast):

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# aigp
```

Related Commands

| Command | Description |
|---------------------------------|-------------------------------------|
| aigp send-cost-community | Sends AiGP value in cost community. |

aigp send-cost-community

To send Accumulated Interior Gateway Protocol (AiGP) value in cost community, use the **aigp send-cost-community** command in appropriate configuration mode. To disable sending AiGP value in cost community, either use the **no** form of this command or the **disable** keyword.

aigp send-cost-community {*cost-id* | **disable**} **poi** {**igp-cost** | **pre-bestpath**} [**transitive**]
no aigp send-cost-community

| | | |
|---------------------------|---------------------|--|
| Syntax Description | <i>cost-comm-id</i> | Specifies the Cost community ID. The range is 0 to 255. |
| | poi | Point of insertion for bestpath calculation. |
| | igp-cost | Configures that cost community be used after iGP distance to next hop. |
| | pre-bestpath | Configures cost community as first step in best path calculation. |
| | transitive | (Optional) Enables transitive cost community |
| | disable | Disables sending AiGP value in cost community. |

Command Default Sending AiGP value in cost community is disabled

Command Modes Neighbor address family configuration

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Cost community point of insertion can be configured either to be pre-bestpath or after igp cost. The **transitive** keyword is not required for iBGP sessions. However, the **transitive** keyword is required for eBGP sessions to convert AiGP metric into cost-community and advertise to the eBGP neighbors.

| | | |
|----------------|----------------|------------------|
| Task ID | Task ID | Operation |
| | bgp | read, write |

Examples The following example shows how to enable sending AiGP value in cost community ID 254 under neighbor address family (IPv4 unicast):

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# aigp send-cost-community 254
```


Related Commands

| Command | Description |
|----------------------------------|--|
| aigp, on page 14 | Enables sending and receiving of accumulated interior gateway protocol (AiGP) attribute. |

allocate-label

To allocate Multiprotocol Label Switching (MPLS) labels for specific IPv4 unicast routes so that the BGP router can send labels with BGP routes to a neighboring router configured for labeledunicast sessions, use the **allocate-label** command in the appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

allocate-label {**route-policy** *route-policy-name* | **all**}
no allocate-label {**route-policy** *route-policy-name* | **all**}

| | | |
|--|---|--|
| Syntax Description | all | Allocates labels for all prefixes |
| | route-policy <i>route-policy-name</i> | Uses a route policy to select prefixes for label allocation. |
| Command Default | No default behavior or values | |
| Command Modes | IPv4 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the allocate-label command with a route policy to trigger BGP to allocate labels for all or a filtered set of global routes (as dictated by the route policy). | |
|  Note | The allocate-label all command is functionally equivalent to the allocate-label route-policy route-policy-name command when the route policy is a pass-all policy. | |
| | | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to enable allocating labels for IPv4 routes:

```
RP/0/RP0/CPU0:router(config)# router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# address family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# allocate-label route-policy policy_A
```

allow vpn default-originate

To configure the router to be enabled to advertise a default route to a configured BGP VPN neighbor, use the **allow vpn default-originate** command in the BGP VRF Address-Family configuration mode. To undo this configuration, use the **no** form of this command.

allow vpn default-originate
no allow vpn default-originate

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | The router cannot advertise a default route to its BGP VPN neighbors. | |
| Command Modes | BGP VRF Address-Family configuration mode | |
| Command History | Release | Modification |
| | Release 4.3.2 | This command was introduced. |
| Usage Guidelines | This command only enables the router to advertise itself as the next-hop router for a default route to its BGP VPN neighbors. To actually forward the default route to a BGP VPN neighbor, you need to run the default-originate command under the BGP neighbor Address-Family configuration mode. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

Example

The following example configuration shows how to enable a BGP router to advertise a default route to its BGP VPN neighbors.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# vrf foo
RP/0/RP0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-vrf-af)# allow vpn default-originate
```

allowas-in

To allow an AS path with the provider edge (PE) autonomous system number (ASN) a specified number of times, use the **allowas-in** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

allowas-in [*as-occurrence-number*]
no allowas-in [*as-occurrence-number*]

| | | |
|--------------------|---|------------------------------|
| Syntax Description | as-occurrence-number (Optional) Number of times a PE ASN is allowed. Range is 1 to 10. | |
| Command Default | No default behavior or values | |
| Command Modes | Address family group configuration Neighbor address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

as-format

To configure the router's Autonomous system number (ASN) notation to asdot format, use the `as-format` command in XR Config mode. To restore the system to its default condition, use the **no** form of this command.

as-format asdot
no

Syntax Description

asdot Specifies the Autonomous system number (ASN) notation to asdot format.

Command Default

The default value, if the `as-format` command is not configured, is `asplain`.

Command Modes

XR Config

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure the ASN notation to the asdot format:

```
RP/0/RP0/CPU0:router(config)# as-format asdot
```

as-path-loopcheck out disable

To disable AS PATH loop checking for outbound updates, use the **as-path-loopcheck out disable** command in an appropriate address family configuration mode. To re-enable the default AS PATH loop checking, use the **no** form of this command.

as-path-loopcheck out disable
no as-path-loopcheck out disable

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | AS PATH loop checking for outbound updates is enabled if there is only one neighbor and disabled if there are multiple neighbors in the update group. | |
| Command Modes | IPv4 address family IPv6 address family | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Configure the as-path-loopcheck out disable command to disable the default behavior of PE router not announcing BGP routes to the CE router if the routes contain an AS number matching the AS number of the receiving CE router.</p> | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to configure **as-path-loopcheck out disable** under IPv6 unicast address family:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#as-path-loopcheck out disable
```

attribute-filter group

To configure attribute-filter group command mode, use the attribute-filter group command in an appropriate configuration mode. To disable attribute-filter group command mode, use the no form of this command.

attribute-filter group *group-name*
no attribute-filter group *group-name*

| Syntax Description | <i>group-name</i> Specifies the name of the attribute-filter group. | | | | |
|---------------------------|---|---------|--------------|---------------|------------------------------|
| Command Default | Attribute-filter group command mode is disabled. | | | | |
| Command Modes | Router configuration Neighbor configuration | | | | |
| Command History | <table> <tr> <th>Release</th><th>Modification</th></tr> <tr> <td>Release 5.0.0</td><td>This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the attribute-filter group command in neighbor configuration mode to configure a specific attribute filter group for a BGP neighbor.</p> | | | | |
| Task ID | <table> <tr> <th>Task ID</th><th>Operation</th></tr> <tr> <td>bgp</td><td>read, write</td></tr> </table> | Task ID | Operation | bgp | read, write |
| Task ID | Operation | | | | |
| bgp | read, write | | | | |

This example shows how to configure the attribute-filter group command mode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#attribute-filter group ag_discard_med
RP/0/RP0/CPU0:router(config-bgp-attrfg)#
```

This example shows how to configure the attribute filter group for a BGP neighbor:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.0.1.101
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 6461
RP/0/RP0/CPU0:router(config-bgp-nbr)#update in filtering
RP/0/RP0/CPU0:router(config-nbr-upd-filter)#attribute-filter group ag_discard_med
```

attribute large-community

To specify the large community BGP path attributes in the created attribute filter and the action to be taken on the update message received from the BGP neighbor, use the **attribute large-community** command in the BGP attribute-filter group command mode.

attribute large-community { discard | treat-as-withdraw }

| | | |
|---------------------------|---|---|
| Syntax Description | discard | Specifies that the large-community BGP path attribute received in a BGP update message from the neighbor must be discarded. |
| | treat-as-withdraw | Indicates that the entire BGP update message that contains the large-community attribute has to be treated as a 'withdraw' message for the Network Layer Reachability Information (NLRI) contained in the update message. |
| Command Default | None | |
| Command Modes | BGP attribute-filter group mode | |
| Command History | Release | Modification |
| | Release 6.3.1 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

The following example creates the attribute-filter named dogbert and specifies the large-community attribute and the action to discard. The dogbert attribute-filter is applied to the BGP neighbor 10.10.10.3. If the large-community BGP path attribute is received in a BGP update message from the neighbor 10.10.10.3, then the attribute will be discarded before further processing of the message.

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 3
RP/0/RP0/CPU0:router(config-bgp)#attribute-filter group dogbert
RP/0/RP0/CPU0:router(config-bgp-attrfg)#attribute large-community discard
RP/0/RP0/CPU0:router(config-bgp-attrfg)#neighbor 10.10.10.3
RP/0/RP0/CPU0:router(config-bgp-nbr)#update in filtering
RP/0/RP0/CPU0:router(config-nbr-upd-filter)#attribute-filter group dogbert
```

bfd (BGP)

To specify a bidirectional forwarding detection (BFD) **multiplier** and **minimum-interval** arguments per neighbor, use the **bfd** command in neighbor address family independent configuration mode. To return to the system defaults, use the **no** form of this command.

Previous to this enhancement, BFD could be configured only in global scope in BGP. This change makes available two new command-line arguments under neighbor address family independent configuration:

```
bfd {multiplier | minimum-interval} value
no bfd {multiplier | minimum-interval} value
```

Syntax Description

| | |
|--------------------------------------|--|
| multiplier <i>value</i> | Specifies the BFD session's multiplier value for the neighbor. |
| minimum-interval <i>value</i> | Specifies the BFD session's minimum-interval value for the neighbor. |

Command Default

No default per neighbor parameters are set.

Command Modes

Neighbor address family independent configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | No modification. |
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the minimum interval is changed using the **bfd minimum-interval** command, the new parameter updates all affected BFD sessions under the command mode in which the minimum interval was changed.

If the multiplier is changed using the **bfd multiplier** command, the new parameter is used to update only the BFD sessions associated with the affected neighbor gets affected.

The assumption is that when BFD fast-detect is enabled under neighbor address family independent configuration, the values for the **multiplier** and **minimum-interval** values are always derived from the per-neighbor values if they are configured; otherwise, they are to be taken from the global BGP configuration mode. In the event that this has not been explicitly stated, then these values are taken to be the default values. Also, the **bfd** arguments can be configured under neighbor-group and session-group and the inheritance adheres to the standard way of BGP configuration inheritance.

Accordingly, there are four cases in which bfd-fast detect is enabled.

This is shown in table below where the BFD value is either multiplier or minimum-interval. Local indicates per NBR value, global is the BGP global value.

| BFD value (global) | BFD value (local) | Result |
|--------------------|-------------------|-------------------|
| Yes | Yes | BFD value (local) |

| BFD value (global) | BFD value (local) | Result |
|-----------------------|----------------------|---------------------|
| Yes | No | BFD value (global) |
| No | Yes | BFD value (local) |
| No | No | BFD value (default) |

bgp as-path-loopcheck

To enable loop checking in the autonomous system path of the prefixes advertised by internal Border Gateway Protocol (iBGP) peers, use the **bgp as-path-loopcheck** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

bgp as-path-loopcheck
no bgp as-path-loopcheck

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | When you do not specify this command, loop checking is performed only for external peers. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>The following example shows how to configure an autonomous system path for loop checking iBGP peers:</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 6 RP/0/RP0/CPU0:router(config-bgp)# bgp as-path-loopcheck</pre> | |

bgp attribute-download

To enable Border Gateway Protocol (BGP) attribute download, use the **bgp attribute-download** command in an appropriate configuration mode. To disable BGP attribute download, use the **no** form of this command.

bgp attribute-download
no bgp attribute-download

Syntax Description This command has no keywords or arguments.

Command Default BGP attribute download is not enabled.

Command Modes IPv4 unicast address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When BGP attribute download is enabled using the **bgp attribute-download** command, BGP reinstalls all routes whose attributes are not currently in the RIB. Likewise, if the user disables BGP attribute download using the no form of the command, BGP reinstalls previously installed routes with a null key, and removes the attributes from the RIB.

Use the **bgp attribute-download** command to enable the Netflow BGP data export function. When attribute download is enabled, BGP downloads the attribute information for prefixes (community, extended community, and as-path) to the Routing Information Base (RIB) and Forwarding Information Base (FIB). This enables FIB to associate the prefixes with attributes and send the Netflow statistics along with the associated attributes.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows the BGP routes before and after BGP attribute download is enabled and shows how to enable BGP attribute download on BGP router 50:

```
RP/0/RP0/CPU0:router# show route bgp

B    100.0.1.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.2.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.3.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.4.0/24 [200/0] via 10.0.101.1, 00:00:37
B    100.0.5.0/24 [200/0] via 10.0.101.1, 00:00:37

RP/0/RP0/CPU0:router(config)# router bgp 50
```

```
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp attribute-download
!
!
!
RP/0/RP0/CPU0:router# show route bgp

B   100.0.1.0/24 [200/0] via 10.0.101.1, 00:00:01
    Attribute ID 0x2
B   100.0.2.0/24 [200/0] via 10.0.101.1, 00:00:01
    Attribute ID 0x2
B   100.0.3.0/24 [200/0] via 10.0.101.1, 00:00:01
    Attribute ID 0x2
B   100.0.4.0/24 [200/0] via 10.0.101.1, 00:00:01
    Attribute ID 0x2
B   100.0.5.0/24 [200/0] via 10.0.101.1, 00:00:01
    Attribute ID 0x2
```

bgp auto-policy-soft-reset disable

To disable an automatic soft reset of Border Gateway Protocol (BGP) peers when their configured route policy is modified, use the **bgp auto-policy-soft-reset disable** command in an appropriate configuration mode. To re-enable automatic soft reset of BGP peers, use the **no** form of this command.

bgp auto-policy-soft-reset disable
no bgp auto-policy-soft-reset disable

| | |
|---------------------------|--|
| Syntax Description | This command has no keywords or arguments. |
|---------------------------|--|

| | |
|------------------------|---|
| Command Default | Automatic soft reset of peers is enabled. |
|------------------------|---|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|



| | |
|-------------|--|
| Note | If the inbound policy changes, it is not always possible to perform a soft reset. This is the case if the neighbor does not support route refresh and soft-reconfiguration inbound is not configured for the neighbor. In such instances, a message is logged in the system log indicating that a manual hard reset is needed. |
|-------------|--|

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

| | |
|-----------------|---|
| Examples | The following example shows how to disable an automatic soft reset of BGP peers when their configured route policy is modified: |
|-----------------|---|

```
RP/0/RP0/CPU0:router(config)# router bgp 6  
RP/0/RP0/CPU0:router(config-bgp)# bgp auto-policy-soft-reset disable
```

bgp bestpath as-path ignore

To ignore the autonomous system path length when calculating preferred paths, use the **bgp bestpath as-path ignore** command in an appropriate configuration mode. To return the software to the default state in which it considers the autonomous system path length when calculating preferred paths, use the **no** form of this command.

bgp bestpath as-path ignore
no bgp bestpath as-path ignore

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | The autonomous system path length is used (not ignored) when a best path is selected. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp bestpath as-path ignore** command to ignore the length of autonomous system paths when the software selects a preferred path. When the best path is selected, if this command is specified, all steps are performed as usual except comparison of the autonomous path length between candidate paths.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to configure the software to ignore the autonomous system length when performing best-path selection:

```
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath as-path ignore
```

| Related Commands | Command | Description |
|------------------|--|--|
| | bgp bestpath compare-routerid , on page 32 | Compares identical routes received from eBGP peers during the best-path selection process and selects the route with the lowest router ID. |

| Command | Description |
|---|---|
| bgp bestpath med always, on page 34 | Allows the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. |
| bgp bestpath med confed, on page 36 | Enables MED comparison among paths learned from confederation peers. |
| bgp bestpath med missing-as-worst, on page 38 | Enables the software to consider a missing MED attribute in a path as having a value of infinity. |

bgp bestpath compare-routerid

To compare identical routes received from external BGP (eBGP) peers during the best-path selection process and select the route with the lowest router ID, use the **bgp bestpath compare-routerid** command in an appropriate configuration mode. To disable comparing identical routes received from eBGP peers during best-path selection, use the **no** form of this command.

bgp bestpath compare-routerid
no bgp bestpath compare-routerid

Syntax Description

This command has no keywords or arguments.

Command Default

The software does not select a new best path if it is the same as the current best path (according to the BGP selection algorithm) except for the router ID.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp bestpath compare-routerid** command to affect how the software selects the best path, in the case where there are two paths of equal cost according to the BGP selection algorithm. This command is used to force the software to select the path with the lower router ID as the best path. If this command is not used, the software continues to use whichever path is currently the best path, regardless of which has the lower router ID.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure the BGP speaker in autonomous system 500 to compare the router IDs of similar paths:

```
RP/0/RP0/CPU0:router(config)# router bgp 500
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath compare-routerid
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp bestpath cost-community ignore

To configure a router that is running the Border Gateway Protocol (BGP) to not evaluate the cost community attribute during the best-path selection process, use the **bgp bestpath cost-community ignore** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

bgp bestpath cost-community ignore
no bgp bestpath cost-community ignore

Syntax Description

This command has no keywords or arguments.

Command Default

The behavior of this command is enabled by default until the cost community attribute is manually configured.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp bestpath cost-community ignore** command to disable the evaluation of the cost community attribute to help isolate problems and troubleshoot issues that relate to BGP path selection. This command can also be used to delay the activation of cost community attribute evaluation so that cost community filtering can be deployed in a large network at the same time.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure a router to not evaluate the cost community attribute during the best-path selection process:

```
RP/0/RP0/CPU0:router(config)# router bgp 500
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath cost-community ignore
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp bestpath med always

To allow the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems, use the **bgp bestpath med always** command in an appropriate configuration mode. To disable considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med always
no bgp bestpath med always

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | The software does not compare MEDs for paths from neighbors in different autonomous systems. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | The MED is one of the parameters that is considered by the software when selecting the best path among many alternative paths. The software chooses the path with the lowest MED. | |
| | By default, during the best-path selection process, the software makes a MED comparison only among paths from the same autonomous system. This command changes the default behavior of the software by allowing comparison of MEDs among paths regardless of the autonomous system from which the paths are received. | |
| | When the bgp bestpath med always command is not enabled and distributed BGP is configured, speakers calculate partial best paths only (executes the best-path steps up to the MED comparison) and send them to BGP Routing Information Base (bRIB). bRIB calculates the final best path (executes all the steps in the best-path calculation). When the bgp bestpath med always command is enabled and distributed BGP is configured, speakers can compare the MED across all ASs, allowing the speaker to calculate a single best path to send it to bRIB. bRIB is the ultimate process that calculates the final best path, but when the bgp bestpath med always command is enabled, the speakers send a single best path instead of potentially sending multiple, partial best paths | |

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to configure the Border Gateway Protocol (BGP) speaker in autonomous system 100 to compare MEDs among alternative paths, regardless of the autonomous system from which the paths are received:

```
RP/0/RP0/CPU0:router(config)# router bgp 100  
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med always
```

Related Commands

| Command | Description |
|---|---|
| bgp bestpath med confed, on page 36 | Enables MED comparison among paths learned from confederation peers. |
| bgp bestpath med missing-as-worst, on page 38 | Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp bestpath med confed

To enable Multi Exit Discriminator (MED) comparison among paths learned from confederation peers, use the **bgp bestpath med confed** command in an appropriate configuration mode. To disable the software from considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med confed
no bgp bestpath med confed

Syntax Description

This command has no keywords or arguments.

Command Default

The software does not compare the MED of paths containing only confederation segments, or paths containing confederation segments followed by an AS_SET, with the MED of any other paths.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, the MED of the following paths is not compared with the MED of any other path:

- Paths with an empty autonomous system path
- Paths beginning with an AS_SET
- Paths containing only confederation segments
- Paths containing confederation segments followed by an AS_SET

Use the **bgp bestpath med confed** command to affect how the following types of paths are treated in the BGP best-path algorithm:

- Paths containing only confederation segments
- Paths containing confederation segments followed by an AS_SET

The MED for paths that start with an AS_SEQUENCE or that start with confederation segments followed by an AS_SEQUENCE only is compared with the MED of other paths that share the same first autonomous system number in the autonomous system sequence (the neighbor autonomous system number). This behavior is not affected by the **bgp bestpath med confed** command.

As an example, suppose that autonomous systems 65000, 65001, 65002, and 65004 are part of a confederation, but autonomous system 1 is not. Suppose that for a particular route, the following paths exist:

- Path 1: 65000 65004, med = 2, IGP metric = 20
- Path 2: 65001 65004, med = 3, IGP metric = 10
- Path 3: 65002 1, med = 1, IGP metric = 30

If the **bgp bestpath med confed** command is enabled, the software selects path 1 as the best path because it:

- Has a lower MED than path 2

- Has a lower IGP metric than path 3

The MED is not compared with path 3 because it has an external autonomous system number (that is, an AS_SEQUENCE) in the path. If the **bgp bestpath med confed** command is not enabled, then MED is not compared between any of these paths. Consequently, the software selects path 2 as the best path because it has the lowest IGP metric.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following command shows how to enable Border Gateway Protocol (BGP) software to compare MED values for paths learned from confederation peers:

```
RP/0/RP0/CPU0:router(config)# router bgp 210
RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med confed
```

Related Commands

| Command | Description |
|---|---|
| bgp bestpath med always, on page 34 | Enables MED comparison among paths from neighbors in different autonomous systems. |
| bgp bestpath med missing-as-worst, on page 38 | Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp bestpath med missing-as-worst

To have the software consider a missing Multi Exit Discriminator (MED) attribute in a path as having a value of infinity, making the path without a MED value the least desirable path, use the **bgp bestpath med missing-as-worst** command in an appropriate configuration mode. To disable considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med missing-as-worst
no bgp bestpath med missing-as-worst

| Syntax Description | This command has no keywords or arguments. | | | | |
|---------------------------|---|---------|--------------|---------------|------------------------------|
| Command Default | The software assigns a value of 0 to the missing MED, causing the path with the missing MED attribute to be considered as the best possible MED. | | | | |
| Command Modes | Router configuration | | | | |
| Command History | <table> <tr> <th>Release</th><th>Modification</th></tr> <tr> <td>Release 5.0.0</td><td>This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | | | | |
| Task ID | <table> <tr> <th>Task ID</th><th>Operations</th></tr> <tr> <td>bgp</td><td>read, write</td></tr> </table> | Task ID | Operations | bgp | read, write |
| Task ID | Operations | | | | |
| bgp | read, write | | | | |
| Examples | <p>The following example shows how to direct the Border Gateway Protocol (BGP) software to consider a missing MED attribute in a path as having a value of infinity, making this path the least desirable path:</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 210 RP/0/RP0/CPU0:router(config-bgp)# bgp bestpath med missing-as-worst</pre> | | | | |

| Related Commands | <table> <tr> <th>Command</th><th>Description</th></tr> <tr> <td>bgp bestpath med always, on page 34</td><td>Enables MED comparison among paths from neighbors in different autonomous systems.</td></tr> <tr> <td>bgp bestpath med confed, on page 36</td><td>Enables MED comparison among paths learned from confederation peers.</td></tr> <tr> <td>show bgp, on page 193</td><td>Displays entries in the BGP routing table.</td></tr> </table> | Command | Description | bgp bestpath med always, on page 34 | Enables MED comparison among paths from neighbors in different autonomous systems. | bgp bestpath med confed, on page 36 | Enables MED comparison among paths learned from confederation peers. | show bgp, on page 193 | Displays entries in the BGP routing table. |
|---|---|---------|-------------|---|--|---|--|---------------------------------------|--|
| Command | Description | | | | | | | | |
| bgp bestpath med always, on page 34 | Enables MED comparison among paths from neighbors in different autonomous systems. | | | | | | | | |
| bgp bestpath med confed, on page 36 | Enables MED comparison among paths learned from confederation peers. | | | | | | | | |
| show bgp, on page 193 | Displays entries in the BGP routing table. | | | | | | | | |

bgp bestpath origin-as allow invalid

To permit all paths marked with an 'invalid' origin-as by RPKI to be considered for BGP best path computation, use the **bgp bestpath origin-as allow invalid** command in the router configuration mode. This configuration can also be made in global address family, neighbor, and neighbor address family submodes. To return the device to default operation, use the **no** form of this command.

For router and global address family configuration mode:

bgp bestpath origin-as allow invalid
no bgp bestpath origin-as allow invalid

For neighbour and neighbor address family configuration mode:

bestpath origin-as allow invalid
no bestpath origin-as allow invalid

Syntax Description

This command has no keywords or arguments.

Command Default

By default, prefixes marked with an 'invalid' origin-as are not considered for BGP best path computation when the router is performing origin-as validation.

Command Modes

Router configuration
Address family
Neighbor
Neighbor Address family

Command History

| Release | Modification |
|---------------|-----------------------------|
| Release 4.2.1 | This command was introduced |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Configuring the **bgp bestpath origin-as allow invalid** command allows paths marked with an 'invalid' origin-as to be considered for best path computation. This can be limited to an address family by configuring it at the address-family submode.

This configuration takes effect only when the **bgp bestpath origin-as use validity** configuration is enabled.

Task ID

| Task ID | Operation |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to permit all invalid paths to be considered for BGP best-path selection in the global mode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#bgp bestpath origin-as allow invalid
```

The following example shows how to permit all invalid paths to be considered for BGP best-path selection in the address family submode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#bgp bestpath origin-as allow invalid
```

The following example shows how to permit all invalid paths to be considered for best-path selection in the neighbor submode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 1.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#bestpath origin-as allow invalid
```

The following example shows how to permit all invalid paths to be considered for best-path selection in the neighbour address-family submode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 1.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#bestpath origin-as allow invalid
```


bgp bestpath origin-as use validity

To enable the BGP Origin AS Validation feature (RPKI) and allow the validity states of BGP paths to be taken into consideration in the bestpath process, use the **bgp bestpath origin-as use validity** command. This can be configured in router configuration mode and address family submode. To return the device to default operation, use the **no** form of this command.

bgp bestpath origin-as use validity
no bgp bestpath origin-as use validity

Syntax Description

This command has no keywords or arguments.

Command Default

By default, the best path computation does not take RPKI states into account.

Command Modes

Router configuration

Address family configuration

Command History

| Release | Modification |
|---------------|-----------------------------|
| Release 4.2.1 | This command was introduced |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

There are three RPKI states - valid, invalid, and not found. When the **bgp bestpath origin-as use validity** command is configured, only paths marked with 'valid' or 'not found' are considered as best path candidates. When the **bgp bestpath origin-as allow invalid** command is configured, paths marked as 'invalid' are also considered but preference is given to routes marked 'valid' over those marked 'invalid'.

Task ID

| Task ID | Operation |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to enable the validity states of BGP paths to affect the path's preference when performing best-path selection:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 50000
RP/0/RP0/CPU0:router(config-bgp)#bgp bestpath origin-as use validity
```

bgp bestpath aigp ignore

To configure a device that is running the Border Gateway Protocol (BGP) to not evaluate the accumulated interior gateway protocol (AIGP) metric during the best path selection process between two paths when one path does not have the AIGP metric, use the **bgp bestpath aigp ignore** command in router configuration mode. To return the device to default operation, use the **no** form of this command.

bgp bestpath aigp ignore
no bgp bestpath aigp ignore

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | AIGP is enabled by default. If this command is not configured, then the accumulated interior gateway protocol (AIGP) metric is evaluated (not ignored) during the best path selection. | |
| Command Modes | Router configuration VRF configuration | |
| Command History | Release | Modification |
| | Release 4.1.1 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>By default, BGP always prefers a path with the AIGP metric. When there are two paths, one with the AIGP metric and the other without, then executing the bgp bestpath aigp ignore command results in BGP performing best path computation as if neither paths has the AIGP metric.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>The following example shows how to configure the software to ignore the accumulated interior gateway protocol (AIGP) metric when performing best-path selection:</p> <pre>RP/0/RP0/CPU0:router#configure RP/0/RP0/CPU0:router(config)#router bgp 50000 RP/0/RP0/CPU0:router(config-bgp)#bgp bestpath aigp ignore</pre> | |

bgp bestpath as-path multipath-relax

To configure a Border Gateway Protocol (BGP) routing process to consider the different autonomous system (AS) paths and load balance multiple paths during best path route selection, use the **bgp bestpath as-path multipath-relax** command. To return the BGP routing process to the default operation, use the **no** form of this command.

bgp bestpath as-path multipath-relax
no bgp bestpath as-path multipath-relax

| | |
|---------------------------|--|
| Syntax Description | This command has no keywords or arguments. |
|---------------------------|--|

| | |
|------------------------|------|
| Command Default | None |
|------------------------|------|

| | |
|----------------------|---|
| Command Modes | Router BGP configuration VRF configuration |
|----------------------|---|

| Command History | <table border="1"> <tr> <th>Release</th> <th>Modification</th> </tr> <tr> <td>Release 3.7.2</td> <td>This command was introduced.</td> </tr> </table> | Release | Modification | Release 3.7.2 | This command was introduced. |
|------------------------|---|---------|--------------|---------------|------------------------------|
| Release | Modification | | | | |
| Release 3.7.2 | This command was introduced. | | | | |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

When BGP multi-pathing is enabled, BGP load-balances user traffic within a single autonomous system (AS). The criteria are that all attributes must match (weight, AS path, etc). However when a device is multi-homed to multiple autonomous systems, BGP cannot load balance traffic between them by default. In order to enable load-balancing of traffic among the multi-homed autonomous systems, the **bgp bestpath as-path multipath-relax** command needs to be enabled. The criteria required for this is that the AS-path length should be equal.

Before you use this command, ensure that BGP is enabled

| Task ID | <table border="1"> <tr> <th>Task ID</th> <th>Operation</th> </tr> <tr> <td>bgp</td> <td>read, write</td> </tr> </table> | Task ID | Operation | bgp | read, write |
|----------------|---|---------|-----------|-----|----------------|
| Task ID | Operation | | | | |
| bgp | read, write | | | | |

Examples

This example shows how to configure multipath load sharing on paths from different autonomous systems in router mode:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 120
RP/0/RP0/CPU0:router(config-bgp)#bgp bestpath as-path multipath-relax
```

bgp client-to-client reflection disable

To disable reflection of routes between route-reflection clients using a Border Gateway Protocol (BGP) route reflector, use the **bgp client-to-client reflection disable** command in address family configuration mode. To re-enable client-to-client reflection, use the **no** form of this command.

bgp client-to-client reflection [*cluster-id cluster-id*] **disable**
no bgp client-to-client reflection [*cluster-id cluster-id*] **disable**

| | |
|---------------------------|--|
| Syntax Description | cluster-id cluster-id (Optional) Cluster ID for which intra-cluster route reflection is to be disabled; maximum of 4 bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295. |
|---------------------------|--|

| | |
|------------------------|---|
| Command Default | Client-to-client reflection is enabled. |
|------------------------|---|

| | |
|----------------------|------------------------------|
| Command Modes | Address family configuration |
|----------------------|------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

By default, the clients of a route reflector that are part of the same cluster are not required to be fully meshed and the routes from a client are reflected to other clients. However, if the clients are fully meshed, route reflection is not required. If the cluster-id is not specified, then this command disables intra-cluster route reflection for all clusters.

Examples

In this example, the three neighbors are fully meshed, so client-to-client reflection is disabled:

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# bgp cluster-id 2
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp client-to-client reflection cluster-id 2 disable
RP/0/RP0/CPU0:router(config-bgp-af)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# bgp cluster-id 2
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# route-reflector-client
RP/0/RP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RP0/CPU0:router(config-bgp-nbrgrp)# exit

RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.21 use neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.22 use neighbor-group rrclients
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.253.23 use neighbor-group rrclients
```

Related Commands

| Command | Description |
|---|---|
| bgp cluster-id, on page 46 | Configures the cluster ID if the BGP cluster has more than one route reflector. |
| route-reflector-client, on page 180 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp cluster-id

To configure the cluster ID if the Border Gateway Protocol (BGP) cluster has more than one route reflector, use the **bgp cluster-id** command in an appropriate configuration mode. To remove the cluster ID, use the **no** form of this command.

bgp cluster-id *cluster-id*
no bgp cluster-id [*cluster-id*]

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | cluster-id Cluster ID of this router acting as a route reflector; maximum of 4 bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295. | |
| Command Default | A cluster ID is not configured. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Together, a route reflector and its clients form a <i>cluster</i>. A cluster of clients usually has a single route reflector. In such instances, the cluster is identified by the software as the router ID of the route reflector. To increase redundancy and avoid a single point of failure in the network, a cluster might have more than one route reflector. If it does, all route reflectors in the cluster must be configured with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>The following example shows how to configure the local router as one of the route reflectors serving the cluster. Neighbor 192.168.70.24 is assigned to the default cluster with cluster-id 1.</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 65534 RP/0/RP0/CPU0:router(config-bgp)# bgp cluster-id 1 RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.24 RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534 RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client</pre> | |

Related Commands

| Command | Description |
|---|---|
| cluster-id, on page 104 | Configures the cluster to which a neighbor belongs. |
| route-reflector-client, on page 180 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp confederation identifier

To specify a Border Gateway Protocol (BGP) confederation identifier, use the **bgp confederation identifier** command in an appropriate configuration mode. To remove the confederation identifier, use the **no** form of this command.

bgp confederation identifier *as-number*
no bgp confederation identifier [*as-number*]

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | as-number Autonomous system (AS) number that internally includes multiple autonomous systems. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. | |
| Command Default | No confederation identifier is configured. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | One way to reduce the internal BGP (iBGP) mesh is to divide an autonomous system into multiple autonomous systems and group them into a single confederation. Each autonomous system is fully meshed within itself, and has a few connections to another autonomous system in the same confederation. Although the peers in different autonomous systems have external BGP (eBGP) sessions, they exchange routing information as if they are iBGP peers. Specifically, the confederation maintains the next hop and local preference information, and that allows you to retain a single Interior Gateway Protocol (IGP) for all autonomous systems. To the outside world, the confederation looks like a single autonomous system. | |
| Task ID | Use the bgp confederation identifier command to specify the autonomous system number for the confederation. This autonomous system number is used when BGP sessions are established with external peers in autonomous systems that are not part of the confederation. | |
| | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to divide the autonomous system into autonomous systems 4001, 4002, 4003, 4004, 4005, 4006, and 4007 with the confederation identifier 5. Neighbor 10.2.3.4 is a router inside the confederation. Neighbor 172.20.16.6 is outside the routing domain confederation. To the outside world, there appears to be a single autonomous system with the number 5.

```
RP/0/RP0/CPU0:router(config)# router bgp 4001
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation identifier 5
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4002
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4003
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4004
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4005
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4006
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 4007
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 4002
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# exit
RP/0/RP0/CPU0:router(config-bgp-nbr)# neighbor 172.20.16.6
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 4009
```

Related Commands

| Command | Description |
|---|---|
| bgp confederation peers, on page 50 | Configures the autonomous systems that belong to the confederation. |

bgp confederation peers

To configure the autonomous systems that belong to the confederation, use the **bgp confederation peers** command in an appropriate configuration mode. To remove the autonomous system from the confederation, use the **no** form of this command.

bgp confederation peers [*as-number*]
no bgp confederation peers [*as-number*]

| Syntax Description | <i>as-number</i> Autonomous system (AS) numbers for Border Gateway Protocol (BGP) peers that belong to the confederation. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. | | | | | | |
|---------------------------|---|------------------------------|---------|---------|------------|--|-----|
| Command Default | No BGP peers are identified as belonging to the confederation. | | | | | | |
| Command Modes | Router configuration | | | | | | |
| Command History | Release | Modification | | | | | |
| | Release 5.0.0 | This command was introduced. | | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>The autonomous systems specified in this command are visible internally to a confederation. Each autonomous system is fully meshed within itself. The bgp confederation identifier, on page 48 command specifies the confederation to which the autonomous systems belong.</p> <p>To specify multiple autonomous systems, enter BGP confederation peer configuration mode then enter one <i>autonomous-system-number</i> for each command line.</p> | | | | | | |
| | <table> <tr> <th>Task ID</th><th>Task ID</th><th>Operations</th></tr> <tr> <td></td><td>bgp</td><td>read, write</td></tr> </table> | | Task ID | Task ID | Operations | | bgp |
| Task ID | Task ID | Operations | | | | | |
| | bgp | read, write | | | | | |
| Examples | <p>The following example shows that autonomous systems 1090 and 1093 belong to a single confederation:</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 1090 RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers 1093</pre> | | | | | | |

The following example shows that autonomous systems 1095, 1096, 1097, and 1098 belong to a single confederation:

```
RP/0/RP0/CPU0:router(config)# router bgp 1095  
RP/0/RP0/CPU0:router(config-bgp)# bgp confederation peers  
RP/0/RP0/CPU0:router(config-bgp-confed-peers)# 1096  
RP/0/RP0/CPU0:router(config-bgp-confed-peers)# 1097  
RP/0/RP0/CPU0:router(config-bgp-confed-peers)# 1098
```

Related Commands

| Command | Description |
|--|---|
| bgp confederation identifier, on page 48 | Specifies a BGP confederation identifier. |

bgp dampening

To enable Border Gateway Protocol (BGP) route dampening or change various BGP route dampening factors, use the **bgp dampening** command in an appropriate configuration mode. To disable route dampening and reset default values, use the **no** form of this command.

bgp dampening [{*half-life* [*reuse* *suppress* *max-suppress-time*]} | **route-policy** *route-policy-name*}]
no bgp dampening [{*half-life* [*reuse* *suppress* *max-suppress-time*]} | **route-policy** *route-policy-name*}]

| | | |
|---------------------------|---|--|
| Syntax Description | half-life | (Optional) Time (in minutes) after which a penalty is decreased. Once the route has been assigned a penalty, the penalty is decreased by half after the half-life period (which is 15 minutes by default). Penalty reduction happens every 5 seconds. Range of the half-life period is from 1 to 45 minutes. |
| | reuse | (Optional) Value for route reuse if the flapping route penalty decreases and falls below the reuse value. When this happens, the route is unsuppressed. The process of unsuppressing routes occurs at 10-second increments. Range is 1 to 20000. |
| | suppress | (Optional) Maximum penalty value. Suppress a route when its penalty exceeds the value specified. When this happens, the route is suppressed. Range is 1 to 20000. |
| | max-suppress-time | (Optional) Maximum time (in minutes) a route can be suppressed. Range is 1 to 255. If the <i>half-life</i> value is allowed to default, the maximum suppress time defaults to 60 minutes. |
| | route-policy <i>route-policy-name</i> | (Optional) Specifies the route policy to use to set dampening parameters. |
| Command Default | Route dampening is disabled. <i>half-life</i> : 15 minutes <i>reuse</i> : 750 <i>suppress</i> : 2000 <i>max-suppress-time</i> : four times <i>half-life</i> value | |
| Command Modes | IPv4 address family configuration IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |

Use the **bgp dampening** command without arguments to enable BGP route dampening with the default parameters. The parameters can be changed by setting them on the command line or specifying them with a routing policy.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to set the *half-life* value to 30 minutes, the *reuse* value to 1500, the *suppress* value to 10000, and the *max-suppress-time* to 120 minutes:

```
RP/0/RP0/CPU0:router(config)# router bgp 50
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp dampening 30 1500 10000 120
```

Related Commands

| Command | Description |
|---|--|
| clear bgp dampening, on page 85 | Clears BGP route dampening information and unsuppresses the suppressed routes. |
| clear bgp flap-statistics, on page 88 | Clears BGP flap statistics. |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |
| show bgp dampened-paths, on page 238 | Displays BGP dampened routes. |
| show bgp flap-statistics, on page 242 | Displays BGP flap statistics. |
| show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |

bgp default local-preference

To change the default local preference value, use the **bgp default local-preference** command in an appropriate configuration mode. To reset the local preference value to the default of 100, use the **no** form of this command.

bgp default local-preference *value*
no bgp default local-preference [*value*]

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | <i>value</i> Local preference value. Range is 0 to 4294967295. Higher values are preferable. | |
| Command Default | Enabled with a value of 100. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Generally, the default value of 100 allows you to easily define a particular path as less preferable than paths with no local preference attribute. The preference is sent to all networking devices in the local autonomous system.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to raise the default local preference value from the default of 100 to 200:

```
RP/0/RP0/CPU0:router(config)# router bgp 200
RP/0/RP0/CPU0:router(config-bgp)# bgp default local-preference 200
```

bgp enforce-first-as disable

To disable the software from enforcing the first autonomous system path (known as the AS path) of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, use the **bgp enforce-first-as disable** command in an appropriate configuration mode. To re-enable enforcing the first AS path of a received route from an eBGP peer to be the same as the remote autonomous system, use the **no** form of this command.

bgp enforce-first-as disable
no bgp enforce-first-as disable

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | By default, the software requires the first autonomous system (in the AS path) of a route received from an eBGP peer to be the same as the remote autonomous system configured. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 3.9.0 | No modification. |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | By default, the software ignores any update received from an eBGP neighbor that does not have the autonomous system configured for that neighbor at the beginning of the AS path. When configured, the command applies to all eBGP peers of the router. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | The following example shows a configuration in which incoming updates from eBGP neighbors are not checked to ensure the first AS number in the AS path is the same as the configured AS number for the neighbor: RP/0/RP0/CPU0:router(config)# router bgp 100 RP/0/RP0/CPU0:router(config-bgp)# bgp enforce-first-as disable | |

Related Commands

| Command | Description |
|---|--|
| enforce-first-as, on page 118 | Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |
| enforce-first-as-disable, on page 120 | Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp fast-external-fallover disable

To disable immediately resetting the Border Gateway Protocol (BGP) sessions of any directly adjacent external peers if the link used to reach them goes down, use the **bgp fast-external-fallover disable** command in an appropriate configuration mode. To disable this function and perform an immediate reset of BGP sessions when a link between peers is lost, use the **no** form of this command.

bgp fast-external-fallover disable
no bgp fast-external-fallover disable

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | disable Disables BGP fast external failover. | |
| Command Default | BGP sessions of any directly adjacent external peers are immediately reset if the link used to reach them goes down. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>By default, BGP sessions of any directly adjacent external peers are immediately reset, which allows the network to recover faster when links go down between BGP peers.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to disable the automatic resetting of BGP sessions:

```
RP/0/RP0/CPU0:router(config)# router bgp 109  
RP/0/RP0/CPU0:router(config-bgp)# bgp fast-external-fallover disable
```

bgp graceful-restart

To enable graceful restart support, use the **bgp graceful-restart** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

bgp graceful-restart
no bgp graceful-restart

Syntax Description This command has no keywords or arguments.

Command Default Graceful restart support is not enabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp graceful-restart** command to enable graceful restart functionality on the router, and also to advertise graceful restart to neighboring routers.



Note The **bgp graceful-restart** command with no options must be used to enable graceful restart before using the **bgp graceful-restart purge-time** , **bgp graceful-restart restart-time** , **bgp graceful-restart stalepath-time** , or **bgp graceful-restart graceful-reset** commands.

When graceful restart is enabled, the BGP graceful restart capability is negotiated with neighbors in the BGP OPEN message when the session is established. If the neighbor also advertises support for graceful restart, then graceful restart is activated for that neighbor session. If the neighbor does not advertise support for graceful restart, then graceful restart is not activated for that neighbor session even though it is enabled locally.

If you enter the **bgp graceful-restart** command after some BGP sessions are established, you must restart those sessions before graceful restart takes effect. Use the **clear bgp** command to restart sessions.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples The following example shows how to enable graceful restart:

```
RP/0/RP0/CPU0:router (config) #router bgp 3
```

RP/0/RP0/CPU0:router(config-bgp)#**bgp graceful-restart**

Related Commands

| Command | Description |
|---|---|
| bgp graceful-restart graceful-reset, on page 60 | Enables a graceful reset if configuration changes force a peer reset. |
| bgp graceful-restart purge-time, on page 61 | Defines the maximum time before stale routes are purged. |
| bgp graceful-restart restart-time, on page 62 | Defines the maximum time advertised to neighbors |
| bgp graceful-restart stalepath-time, on page 63 | Defines the maximum time to wait for the End-of-RIB message from a neighbor that has been restarted before deleting learned routes. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 290 | Displays BGP process information. |

bgp graceful-restart graceful-reset

To invoke a graceful restart when configuration changes force a peer reset, use the **bgp graceful-restart graceful-reset** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

bgp graceful-restart graceful-reset
no bgp graceful-restart graceful-reset

Syntax Description

This command has no keywords or arguments.

Command Default

Graceful restart is not invoked when a configuration change forces a peer reset.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

BGP graceful restart must be enabled using the **bgp graceful-restart** command before enabling graceful reset using the **bgp graceful-restart graceful-reset** command.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to enable graceful reset:

```
RP/0/RP0/CPU0:router(config)#router bgp 3
RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart graceful-reset
```

Related Commands

| Command | Description |
|--|--|
| bgp graceful-restart, on page 58 | Enables a graceful restart. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 290 | Displays BGP process information. |

bgp graceful-restart purge-time

To specify the maximum time before stale routes are purged from the routing information base (RIB) when the local BGP process restarts, use the **bgp graceful-restart purge-time** command in an appropriate configuration mode. To set the purge timer time to its default value, use the **no** form of this command.

bgp graceful-restart purge-time *seconds*
no bgp graceful-restart purge-time *seconds*

| | |
|---------------------------|--|
| Syntax Description | <i>seconds</i> Maximum time before stale routes are purged. Time in seconds. Range is 0 to 6000. |
|---------------------------|--|

| | |
|------------------------|----------------------|
| Command Default | <i>seconds</i> : 600 |
|------------------------|----------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the purge time using the **bgp graceful-restart purge-time** command.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

| | |
|-----------------|--|
| Examples | The following example shows how to change the BGP purge time to 800 seconds: |
|-----------------|--|

```
RP/0/RP0/CPU0:router(config)# router bgp 3
RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart purge-time 800
```

| Related Commands | Command | Description |
|-------------------------|---|--|
| | bgp graceful-restart , on page 58 | Enables a graceful restart. |
| | show bgp , on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors , on page 256 | Displays information about BGP connections to neighbors. |
| | show bgp process , on page 290 | Displays BGP process information. |

bgp graceful-restart restart-time

To specify a user-predicted local BGP process maximum restart time, which is advertised to neighbors during session establishment, use the **bgp graceful-restart restart-time** command in an appropriate configuration mode. To set this restart time to its default value, use the **no** form of this command.

bgp graceful-restart restart-time *seconds*

no bgp graceful-restart restart-time *seconds*

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | <i>seconds</i> Maximum time advertised to neighbors. Time in seconds. Range is 1 to 4095. | |
| Command Default | <i>seconds</i> : 120 | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>BGP graceful restart must be enabled using the bgp graceful-restart command before setting the restart timer using the bgp graceful-restart restart-time command.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>The following example shows how to change the BGP graceful restart time to 400 seconds:</p> <pre>RP/0/RP0/CPU0:router(config)#router bgp 3 RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart restart-time 400</pre> | |

| | | |
|-------------------------|--|--|
| Related Commands | Command | Description |
| | bgp graceful-restart, on page 58 | Enables a graceful restart. |
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |
| | show bgp process, on page 290 | Displays BGP process information. |

bgp graceful-restart stalepath-time

To specify the maximum time to wait for an End-of-RIB message after a neighbor restarts, use the **bgp graceful-restart stalepath-time** command in an appropriate configuration mode. To set the stalepath timer time to its default value, use the **no** form of this command.

bgp graceful-restart stalepath-time *seconds*
no bgp graceful-restart stalepath-time *seconds*

| | |
|---------------------------|--|
| Syntax Description | <i>seconds</i> Maximum wait time. Time in seconds. Range is 1 to 4095. |
|---------------------------|--|

| | |
|------------------------|----------------------|
| Command Default | <i>seconds</i> : 360 |
|------------------------|----------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the stalepath time using the **bgp graceful-restart stalepath-time** command.

If the stalepath time is exceeded before an End-of-RIB message is received from a neighbor, paths learned from the neighbor are purged from the BGP routing table.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to change the stalepath time to 750 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 3
RP/0/RP0/CPU0:router(config-bgp)# bgp graceful-restart stalepath-time 750
```

| | | |
|-------------------------|--|--|
| Related Commands | Command | Description |
| | bgp graceful-restart, on page 58 | Enables a graceful restart. |
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |

| Command | Description |
|---|-----------------------------------|
| show bgp process, on page 290 | Displays BGP process information. |

bgp label-delay

To enable delay for Border Gateway Protocol (BGP) batch label processing, use the **bgp label-delay** command in an appropriate configuration mode. To disable delay in batch import processing, use the no form of this command.

bgp label-delay *seconds milliseconds*
no bgp label-delay

| | | |
|---------------------------|---|--|
| Syntax Description | <i>seconds</i> | Specifies batch label processing delay in seconds. Range is 0 to 10 seconds. |
| | <i>milliseconds</i> | Specifies batch label processing delay in milliseconds. Range is 0 to 999 seconds. |
| Command Default | No delay is configured. | |
| Command Modes | Address-family IPv4 Unicast | |
| | Address-family IPv6 Unicast | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to set delay in batch import processing as two seconds and zero milliseconds:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#bgp label-delay 2 0
```

bgp log neighbor changes disable

To disable logging of Border Gateway Protocol (BGP) neighbor resets, use the **bgp log neighbor changes disable** command in an appropriate configuration mode. To re-enable logging of BGP neighbor resets, use the **no** form of this command.

bgp log neighbor changes disable
no bgp log neighbor changes disable

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | This command has no keywords or arguments. | |
| Command Default | BGP neighbor changes are logged. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Logging of BGP neighbor status changes (up or down) and resets is used for troubleshooting network connectivity problems and measuring network stability. Unexpected neighbor resets might indicate high error rates or high packet loss in the network, and should be investigated. | |
| | Status change message logging does not substantially affect performance, unlike, for example, enabling per-BGP update debugging. If the UNIX syslog facility is enabled, messages are sent by the software to the UNIX host running the syslog daemon so that the messages can be stored and archived on disk. If the UNIX syslog facility is not enabled, the status change messages are kept in the internal buffer of the router, and are not stored to disk. | |
| | The neighbor status change messages are not tracked if the bgp log neighbor changes disable command is disabled, except for the last reset reason, which is always available as output of the show bgp neighbors command. | |
| | Up and down messages for BGP neighbors are logged by the software by default. Use the bgp log neighbor changes disable command to stop logging BGP neighbor changes. | |

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to prevent the logging of neighbor changes for BGP:

```
RP/0/RP0/CPU0:router(config)# router bgp 65530
```

```
RP/0/RP0/CPU0:router(config-bgp)# bgp log neighbor changes disable
```

Related Commands

| Command | Description |
|---|--|
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |

bgp maximum neighbor

To control the maximum number of neighbors that can be configured on the router, use the **bgp maximum neighbor** command in an appropriate configuration mode. To set the neighbor limit to the default value, use the **no** form of this command.

bgp maximum neighbor *limit*
no maximum neighbor [*limit*]

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | <i>limit</i> Maximum number of neighbors. Range is 1 to 15000. | |
| Command Default | Default limit is 4000 | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Any attempt to configure the neighbor limit below 1 or above 15000 fails. Similarly, attempting to configure the limit below the number of neighbors currently configured fails. For example, if there are 3250 neighbors configured, you cannot set the <i>limit</i> below 3250.</p> | |
| Task ID | Task ID | Operations |
| | bgp | write |
| Examples | <p>The following example shows how to change the default maximum neighbor limit and set it to 1200:</p> <pre>RP/0/RP0/CPU0:router(config)#router bgp 65530 RP/0/RP0/CPU0:router(config-bgp)# bgp maximum neighbor 1200</pre> | |

bgp nexthop resolution allow-default

By default, the next hop resolution in BGP does not take the default route into account. By configuring this command, the default route is used for resolving the next-hop of BGP routes. The next hop resolution is important in deciding if the next hop for a BGP route is accessible or not.

If the BGP route has an inaccessible next hop, the route does not have a best path and will not be advertised.



Warning

Applying or removing this command will reset all BGP neighbors.

bgp nexthop resolution allow-default

Syntax Description

allow-default Enable nexthops resolution using default route.

Command Default

This applies to IPv4 and IPv6. The default route is 0.0.0.0/0 for IPv4 and ::/0 for IPv6.

Command Modes

XR Config mode

Command History

| Release | Modification |
|-------------|------------------------------|
| Release 6.2 | This command was introduced. |

Task ID

| Task ID | Operation |
|---------|----------------|
| bgp | read, write |

The following example shows how to configure BGP with `nexthop resolution allow-default` :

```
RP/0/0/CPU0:R1(config)#router bgp 65000
RP/0/0/CPU0:R1(config-bgp)#nexthop resolution allow-default
```

"NEXTHOP configuration changed" is seen as the last reset reason with the `show bgp neighbor` command when the `nexthop resolution allow-default` command is applied or removed:

```
RP/0/0/CPU0:R1#show bgp neighbor 10.0.0.2
...
  Last reset 00:01:59, due to NEXTHOP configuration changed

RP/0/0/CPU0:R1#show bgp neighbor 2001:db8:1::2
...
  Last reset 00:02:47, due to NEXTHOP configuration changed
```

bgp redistribute-internal

To allow the redistribution of internal Border Gateway Protocol (iBGP) routes into an Interior Gateway Protocol (IGP), such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF), use the **bgp redistribute-internal** command in an appropriate configuration mode. To disable the redistribution of iBGP routes into IGPs, use the **no** form of this command.

bgp redistribute-internal
no bgp redistribute-internal

Syntax Description This command has no keywords or arguments.

Command Default By default, iBGP routes are not redistributed into IGPs.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use of the **bgp redistribute-internal** command requires the **clear route *** command to be issued to reinstall all BGP routes into the IP routing table.



Note Redistributing iBGP routes into IGPs may cause routing loops to form within an autonomous system. Use this command with caution.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to redistribute iBGP routes into OSPF:

```
RP/0/RP0/CPU0:router(config)#router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# bgp redistribute-internal
RP/0/RP0/CPU0:router(config-bgp)# exit
RP/0/RP0/CPU0:router(config)# router ospf areal
RP/0/RP0/CPU0:router(config-router)# redistribute bgp 1
RP/0/RP0/CPU0:router(config-router)# end
RP/0/RP0/CPU0:router# clear route *
```

Related Commands

| Command | Description |
|---|---------------------------|
| clear bgp, on page 82 * | Resets all BGP neighbors. |
| clear route * | Resets all routes. |

bgp router-id

To configure a fixed router ID for a Border Gateway Protocol (BGP)-speaking router, use the **bgp router-id** command in an appropriate configuration mode. To disable a fixed router ID, use the **no** form of this command.

bgp router-id *ip-address*
no bgp router-id [*{ip-address}*]

Syntax Description

ip-address IP Version 4 (IPv4) address to use as the router ID. Normally, this should be an IPv4 address assigned to the router.

Command Default

If no router ID is configured in BGP, BGP attempts to use the global router ID if one is configured and available. Otherwise, BGP uses the highest IP address configured on a loopback interface.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you do not use the **bgp router-id** command to configure a router ID, an IP address is not configured on any loopback interface, and no global router ID is configured, BGP neighbors remain down.

For more details on router IDs, see the *Routing Configuration Guide for Cisco NCS 6000 Series Routers*.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure the local router with the router ID of 192.168.70.24:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#bgp router-id 192.168.70.24
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp scan-time

To configure scanning intervals of Border Gateway Protocol (BGP)-speaking networking devices, use the **bgp scan-time** command in an appropriate configuration mode. To restore the scanning interval to its default value, use the **no** form of this command.

bgp scan-time *seconds*
no bgp scan-time
seconds

| | |
|---------------------------|---|
| Syntax Description | <i>seconds</i> Scanning interval (in seconds) of BGP routing information. Range is 5 to 3600 seconds. |
|---------------------------|---|

| | |
|------------------------|--|
| Command Default | The default scanning interval is 60 seconds. |
|------------------------|--|

| | |
|----------------------|--|
| Command Modes | Router configuration IPv4 address family configuration IPv6 address family configuration |
|----------------------|--|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

Use the **bgp scan-time** command to change how frequently the software processes scanner tasks, such as conditional advertisement, dynamic MED changes, and periodic maintenance tasks.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

| | |
|-----------------|--|
| Examples | The following example shows how to set the scanning interval for IPv4 unicast to 20 seconds: |
|-----------------|--|

```
RP/0/RP0/CPU0:router(config)# router bgp 64500
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# bgp scan-time 20
```

| | | |
|-------------------------|---------------------------------------|--|
| Related Commands | Command | Description |
| | show bgp, on page 193 | Displays entries in the BGP routing table. |

bgp update-delay

To set the maximum initial delay for a Border Gateway Protocol (BGP)-speaking router to send the first updates, use the **bgp update-delay** command in an appropriate configuration mode. To restore the initial delay to its default value, use the **no** form of this command.

bgp update-delay *seconds* [**always**]
no bgp update-delay [*seconds*][**always**]

| | | |
|---------------------------|----------------------|---|
| Syntax Description | <i>seconds</i> | Delay in seconds for the router to send the first updates. Range is 0 to 3600. |
| | always | (Optional) Specifies that the router always wait for the update delay time, even if all neighbors have finished sending their initial updates sooner. |
| Command Default | 120 seconds | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When BGP is started, it waits a specified period of time for its neighbors to establish peering sessions and to complete sending their initial updates. After all neighbors complete their initial updates, or after the update delay timer expires, the best path is calculated for each route, and the software starts sending advertisements out to its peers. This behavior improves convergence time. If the software were to advertise a route as soon as it learned it, it would have to readvertise the route each time it learned a new path that was preferred over all previously learned paths.

Use the **bgp update-delay** command to tune the maximum time the software waits after the first neighbor is established until it starts calculating best paths and sending out advertisements.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples

The following example shows how to set the maximum initial delay to 240 seconds:

```
RP/0/RP0/CPU0:router(config)#router bgp 64530
RP/0/RP0/CPU0:router(config-bgp)# bgp update-delay 240
```

capability additional-paths receive

To advertise capability of receiving additional paths to the peer, use the **capability additional-paths receive** command in neighbor or neighbor-group or session-group configuration mode. To disable the capability of receiving additional paths, use the **no** form of this command.

capability additional-paths receive [**disable**]
no **capability additional-paths receive**

| | |
|---------------------------|---|
| Syntax Description | disable Disables advertising capability of receiving additional paths. |
|---------------------------|---|

| | |
|------------------------|-------------------------|
| Command Default | Capability is disabled. |
|------------------------|-------------------------|

| | |
|----------------------|---|
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration |
|----------------------|---|

| Command History | <table><tr><th>Release</th><th>Modification</th></tr><tr><td>Release 5.0.0</td><td>This command was introduced.</td></tr></table> | Release | Modification | Release 5.0.0 | This command was introduced. |
|------------------------|---|---------|--------------|---------------|------------------------------|
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

Use the **capability additional-paths receive** command to selectively enable or disable additional paths receive capability negotiation for a particular neighbor or neighbor-group or session-group.

If you enter the **capability additional-paths receive** command after some BGP sessions are established, you must restart those sessions for the new configuration to take effect. Use the **clear bgp** command to restart sessions.

| Task ID | <table><tr><th>Task ID</th><th>Operation</th></tr><tr><td>bgp</td><td>read, write</td></tr></table> | Task ID | Operation | bgp | read, write |
|----------------|---|---------|-----------|-----|----------------|
| Task ID | Operation | | | | |
| bgp | read, write | | | | |

The following example shows how to advertise capability of receiving additional paths:

```
RP/0/RP0/CPU0:router(config)#router bgp 100  
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.2.3.4  
RP/0/RP0/CPU0:router(config-bgp-nbr)#capability additional-paths receive
```

capability additional-paths send

To advertise capability of sending additional paths to the peer, use the **capability additional-paths send** command in neighbor or neighbor-group or session-group configuration mode. To disable the capability of sending additional paths, use the **no** form of this command.

capability additional paths send [disable]
no capability additional paths send

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | disable Disables advertise additional paths send capability | |
| Command Default | Capability is disabled. | |
| Command Modes | Neighbor configuration | |
| | Neighbor group configuration | |
| | Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the capability additional paths send command to selectively enable or disable additional paths send capability negotiation for a particular neighbor or neighbor-group or session-group. | |
| | You must restart the BGP sessions for the new configuration to take effect. Use the clear bgp command to restart sessions. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

The following example shows how to advertise capability of sending additional paths to the peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)# capability additional-paths send
```

capability orf prefix

To advertise prefix list-based Outbound Route Filter (ORF) capability to the Border Gateway Protocol (BGP) peer, use the **capability orf prefix** command in an appropriate configuration mode. To remove the **capability orf prefix** command from the configuration file and restore the system to its default condition in which the software does not advertise the capability, use the **no** form of this command.

```
capability orf prefix {receive | send | both | none}
no capability orf prefix [{receive | send | both | none}]
```

Syntax Description

| | |
|----------------|--|
| receive | Sets the capability to receive the ORF from a specified neighbor. |
| send | Sets the capability to send the ORF to a specified neighbor. |
| both | Sets the capability to receive and send the ORF from or to a specified neighbor. |
| none | Sets the capability to no for ORF receive or send from or to a specified neighbor. |

Command Default

The routing device does not receive or send route prefix filter lists.

Command Modes

IPv4 address family group configuration
IPv6 address family group configuration
IPv4 neighbor address family configuration
IPv4 neighbor group address family configuration
IPv6 neighbor group address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The advertisement of the prefix list ORF capability by a BGP speaker indicates whether the speaker can send prefix lists to the specified neighbor and whether it accepts prefix lists from the neighbor. The speaker sends a prefix list if it indicated the ability to send them, and if the neighbor indicated it was willing to accept them. Similarly, the neighbor sends a prefix list to the speaker if it indicated the ability to send them and the speaker indicated the willingness to accept them.



Note

The capability orf and prefix list filter specified by orf route-policy must be explicitly configured.

If the neighbor sends a prefix list and the speaker accepts it, the speaker applies the received prefix list, plus any locally configured outbound filters, to limit its outbound routing updates to the neighbor. Increased filtering

prevents unwanted routing updates between neighbors and reduces resource requirements for routing update generation and processing.

Use the **capability orf prefix** command to set whether to advertise send and receive capabilities to the specified neighbor.

**Note**

Sending a receive capability can adversely affect performance, because updates sent to that neighbor cannot be replicated for any other neighbors.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure the **capability orf prefix** command:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# route-policy orfqq
RP/0/RP0/CPU0:router(config-rpl)# if orf prefix in (10.0.0.0/8 ge 20) then
RP/0/RP0/CPU0:router(config-rpl)# pass
RP/0/RP0/CPU0:router(config-rpl)# endif
RP/0/RP0/CPU0:router(config-rpl)# if orf prefix in (1910::16 ge 120) then
RP/0/RP0/CPU0:router(config-rpl)# pass
RP/0/RP0/CPU0:router(config-rpl)# endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config)# router bgp 65530
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy pass-all out
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# capability orf prefix both
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# orf route-policy orfqq
```


Related Commands

| Command | Description |
|---|--|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| show bgp neighbors, on page 256 | Displays information about BGP neighbors. Use the received prefix-filter keywords to display information on the prefix list filter. |

capability suppress

To suppress capabilities from being advertised to the BGP peer, use the **capability suppress** command in the appropriate configuration mode. To remove the **capability suppress** command from the configuration and restore the system to the default condition, in which the software advertises the capability, use the **no** form of this command.

capability suppress {4-byte-as | all | extended-nexthop-encoding} [inheritance-disable]

| | | |
|---|---|--|
| Syntax Description | 4-byte-as | Suppresses advertisement of the 4-byte AS capability to BGP peers. |
| | all | Suppresses advertisement of all capabilities to BGP peers. |
| | inheritance-disable | Prevents capability suppress being inherited from the parent. |
| | | |
| Command Default | 4-byte-as capability is advertised to the BGP peer. | |
| Command Modes | Neighbor configuration | |
| | Neighbor group configuration | |
| | Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | By default, the software advertises the 4-byte AS capability to BGP peers. To override this default behavior, use the capability suppress 4-byte-as command under the command modes listed in the "Command Modes" section. If configured under the neighbor group or session group, all neighbors using the group inherit the configuration. Use the no option to remove the command. | |
| <div></div> <div>Caution</div> | The BGP session resets automatically, if the 4-byte AS capability of an existing BGP session is changed by configuring capability suppress 4-byte-as or capability suppress 4-byte-as inheritance-disable . | |
| | | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | The following example shows how to configure the capability suppress 4-byte-as command: RP/0/RP0/CPU0:router# show bgp nei 10.3.3.3 conf neighbor 10.3.3.3 remote-as 65000 | |

```
update-source Loopback0      [n:internal]
address-family ipv4 unicast [n:internal]
```

```
RP/0/RP0/CPU0:router#show bgp nei 10.3.3.3
```

```
BGP neighbor is 10.3.3.3
Remote AS 65000, local AS 65000, internal link
Description: PE3
Remote router ID 10.3.3.3
BGP state = Established, up for 1w0d
Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
Precedence: internet
Neighbor capabilities:
  Route refresh: advertised and received
  4-byte AS: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 25962 messages, 0 notifications, 0 in queue
Sent 25968 messages, 1 notifications, 0 in queue
Minimum time between advertisement runs is 0 seconds
```

```
For Address Family: IPv4 Unicast
BGP neighbor version 1
Update group: 0.3
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
An EoR was received during read-only mode

Connections established 2; dropped 1
Last reset 1w0d, due to BGP Notification sent: hold time expired
Time since last notification sent to neighbor: 1w0d
Error Code: hold time expired
Notification data sent: None
```

```
RP/0/RP0/CPU0:router(config)#router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.3.3.3
RP/0/RP0/CPU0:router(config-bgp-nbr)#capability suppress 4-byte-as
RP/0/RP0/CPU0:router(config-bgp-nbr)#commit
RP/0/RP0/CPU0:router(config-bgp-nbr)#end
```

```
RP/0/RP0/CPU0:router# show bgp nei 10.3.3.3
```

```
BGP neighbor is 10.3.3.3
Remote AS 65000, local AS 65000, internal link
Description: PE3
Remote router ID 10.3.3.3
BGP state = Established, up for 00:00:16
Last read 00:00:11, hold time is 180, keepalive interval is 60 seconds
Precedence: internet
Neighbor capabilities:
  Route refresh: advertised and received
  Address family IPv4 Unicast: advertised and received
  Capability 4-byte-as suppress is configured
Received 25966 messages, 0 notifications, 0 in queue
Sent 25972 messages, 1 notifications, 0 in queue
Minimum time between advertisement runs is 0 seconds
```

```
For Address Family: IPv4 Unicast
BGP neighbor version 1
Update group: 0.2
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
```


An EoR was received during read-only mode

Connections established 3; dropped 2
Last reset 00:00:43, due to Capabilty 4-byte-as configuration changed
Time since last notification sent to neighbor: 1w0d
Error Code: hold time expired
Notification data sent: None

With the **inheritance-disable** keyword:

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.101.1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# capability suppress 4-byte-as inheritance-disable
```

```
RP/0/RP0/CPU0:router# show bgp neighbor 10.0.101.1 config  
neighbor 10.0.101.1  
  remote-as 1 []  
  address-family ipv4 unicast []
```

```
RP/0/RP0/CPU0:router# show bgp neighbor 10.0.101.1  
BGP neighbor is 10.0.101.1  
  Remote AS 1, local AS 100, external link  
  Remote router ID 0.0.0.0  
  BGP state = Idle  
  Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds  
  Precedence: internet  
  Received 0 messages, 0 notifications, 0 in queue  
  Sent 0 messages, 0 notifications, 0 in queue  
  Minimum time between advertisement runs is 30 seconds
```

clear bgp

To reset a group of Border Gateway Protocol (BGP) neighbors, use the **clear bgp** command in XR EXEC mode.

clear bgp [{**ipv4** {**unicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **all** | **labeled-unicast**} | **all** }]

Syntax Description

| | |
|------------------------|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |

Command Default

No default behavior or values

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear bgp** command to reset the sessions of the specified group of neighbors (hard reset); it removes the TCP connection to the neighbor, removes all routes received from the neighbor from the BGP table, and then re-establishes the session with the neighbor.

If the **graceful** keyword is specified, the routes from the neighbor are not removed from the BGP table immediately, but are marked as stale. After the session is re-established, any stale route that has not been received again from the neighbor is removed.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to hard reset neighbor 10.0.0.1:

```
RP/0/RP0/CPU0:router# clear bgp 10.0.0.1
```

Related Commands

| Command | Description |
|---|--|
| clear bgp self-originated, on page 98 | Clears self-originated routes. |
| clear bgp soft, on page 101 | Soft resets a group of BGP neighbors. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |

cef consistency-hashing auto-recovery

To enable automatic recovery of failed ECMP links and the sessions distributed due the ECMP link failure, use the **cef consistent-hashing auto-recovery** command in global configuration mode.

cef consistent-hashing auto-recovery

| | |
|---------------------------|--|
| Syntax Description | This command has no keywords or arguments. |
|---------------------------|--|

| | |
|------------------------|--|
| Command Default | Failed ECMP links are not automatically recovered. |
|------------------------|--|

| | |
|----------------------|----------------------|
| Command Modes | Global configuration |
|----------------------|----------------------|

| | | |
|------------------------|----------------|-----------------------------|
| Command History | Release | Modification |
| | Release 6.5.1 | The command was introduced. |

| | |
|-------------------------|--|
| Usage Guidelines | Configuring the command does not alter the current state. The command takes effect on the next link down or up events. |
|-------------------------|--|

| | | |
|----------------|----------------|------------------|
| Task ID | Task ID | Operation |
| | ipv4 | read, write |

Example

```
Router# configure
Router(config)# cef consistent-hashing auto-recovery
```

clear bgp dampening

To clear Border Gateway Protocol (BGP) route dampening information and unsuppress the suppressed routes, use the **clear bgp dampening** command in XR EXEC mode.

clear bgp dampening

| | | |
|--------------------|---|--|
| Syntax Description | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | labeled-unicast | Specifies labeled-unicast address prefixes. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | all | For address family, specifies prefixes for all address families. |
| | <i>ip-address</i> | (Optional) IP address of the network about which to clear dampening information. |
| | <i>/mask-length</i> | (Optional) Network mask applied to the IP address. |
| Command Default | If no IP address is specified, dampening information for all routes is cleared. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Always use the clear bgp dampening command for an individual address-family. The all option for address-families with clear bgp dampening should never be used during normal functioning of the system. For example, use <pre>clear bgp ipv4 unicast dampening prefix x.x.x./y</pre> | |
| Task ID | Task ID | Operations |
| | bgp | execute |
| Examples | The following example shows how to clear the route dampening information for all 172.20.0.0/16 IPv4 unicast paths: | |
| | <pre>RP/0/RP0/CPU0:router# clear bgp ipv4 unicast dampening 172.20.0.0/16</pre> | |

Related Commands

| Command | Description |
|--|---|
| bgp dampening, on page 52 | Enables BGP route dampening or changes various BGP route dampening factors. |
| show bgp dampened-paths, on page 238 | Displays BGP dampened routes. |

clear bgp external

To clear all Border Gateway Protocol (BGP) external peers, use the **clear bgp external** command in XR EXEC mode.

clear bgp external

| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
|--------------------|---|--|
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | graceful | (Optional) Clears all external peers with a hard reset and a graceful restart. This option is available when an address family is not specified. |
| Command Default | No default behavior or value | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | execute |
| Examples | The following example shows how to clear all BGP external peers: RP/0/RP0/CPU0:router# clear bgp external | |

clear bgp flap-statistics

To clear Border Gateway Protocol (BGP) flap counts for a specified group of routes, use the **clear bgp flap-statistics** command in XR EXEC mode.

clear bgp flap-statistics

| | | |
|---------------------------|---|--|
| Syntax Description | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | labeled-unicast | Specifies labeled unicast address prefixes. |
| | all | For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | all | For address family, specifies prefixes for all address families. |
| | regexp <i>regexp</i> | (Optional) Clears flap statistics for routes whose AS paths match the regular expression. |
| | route-policy <i>route-policy-name</i> | (Optional) Clears flap statistics for the specific route policy. |
| | <i>network</i> | (Optional) Network for which flap counts are to be cleared. |
| | <i>/mask-length</i> | (Optional) Network mask of the network for which flap counts are to be cleared. |
| | <i>ip-address</i> | (Optional) Neighbor address. Clears only flap statistics for routes received from this neighbor. |
| Command Default | No default behavior or value | |
| Command Modes | EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | execute |

Examples

The following example shows how to clear the flap count for all routes (in all address families) originating in autonomous system 1:

```
RP/0/RP0/CPU0:router#clear bgp all all flap-statistics regexp _1$
```

The following example shows how to clear the flap count for all IPv4 unicast routes received from neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router# clear bgp ipv4 unicast flap-statistics 172.20.1.1
```

clear bgp nexthop performance-statistics

To reset the number of received notifications and the cumulative processing time for the Border Gateway Protocol (BGP) next hop, use the **clear bgp nexthop performance-statistics** command in XR EXEC mode.

clear bgp nexthop performance-statistics

| | | |
|---------------------------|--|--|
| Syntax Description | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | labeled-unicast | Specifies labeled unicast address prefixes. |
| | all | For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | all | For address family, specifies prefixes for all address families. |
| Command Default | No default behavior or values | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | <p>Use the clear bgp nexthop performance-statistics command to reset the total number of notifications received from the Routing Information Base (RIB) and the cumulative next-hop processing time. The following information is cleared from the show bgp nexthops command output:</p> <ul style="list-style-type: none"> • Total critical notifications received • Total noncritical notifications received • Best path deleted after last walk • Best path changed after last walk • Next-hop table total number of critical and noncritical notifications (Notf) and the time of the last notification received from the RIB (LastRIB) columns (only entries that have a status of unreachable [UR]) | |
| Task ID | Task ID | Operations |
| | bgp | execute |

Related Commands

| Command | Description |
|--|--|
| show bgp nexthops, on page 275 | Displays information about the BGP next-hop notifications. |

clear bgp nexthop registration

To reregister a specified next hop with the Routing Information Base (RIB), use the **clear bgp nexthop registration** command in XR EXEC mode.

clear bgp nexthop registration nexthop-address *nexthop-address*

| | | |
|---------------------------|--|--|
| Syntax Description | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | labeled-unicast | Specifies labeled-unicast address prefixes. |
| | all | For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | all | For address family, specifies prefixes for all address families. |
| | <i>nexthop-address</i> | Address of the next hop. |
| Command Default | No default behavior or values | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the clear bgp nexthop registration command to perform an asynchronous registration of the next hop with the RIB. The show bgp nexthops command output shows a critical notification as the LastRIBEvent for the next hop when the clear bgp nexthop registration command is used. | |
| Task ID | Task ID | Operations |
| | bgp | execute |
| Examples | The following example shows how to reregister the next hop with the RIB: | |
| | <pre>RP/0/RP0/CPU0:router# clear bgp nexthop registration 10.1.1.1</pre> | |

Related Commands

| Command | Description |
|--|--|
| show bgp nexthops, on page 275 | Displays information about the BGP next-hop notifications. |

clear bgp peer-drops

To clear the connection-dropped counter, use the **clear bgp peer-drops** command in XR EXEC mode.

clear bgp peer-drops *{*ip-address}*

| | | |
|---------------------------|---|--|
| Syntax Description | * | Specifies all BGP neighbors. |
| | <i>ip-address</i> | IP address of a specific network neighbor. |
| Command Default | No default behavior or values | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | execute |
| Examples | The following example shows how to clear the connection-dropped counter for all BGP neighbors: RP/0/RP0/CPU0:router# clear bgp peer-drops * | |
| Related Commands | Command | Description |
| | show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |

clear bgp performance-statistics

To clear the performance statistics for all address families, use the **clear bgp performance-statistics** command.

clear bgp performance-statistics

| | |
|---------------------------|--|
| Syntax Description | This command has no keywords or arguments. |
|---------------------------|--|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | execute |

| | |
|-----------------|---|
| Examples | The following example shows how to clear the performance statistics for all address families: |
|-----------------|---|

```
RP/0/RP0/CPU0:router# clear bgp performance-statistics
```

clear bgp rpki

To clear resource public key infrastructure (RPKI) cache-server (rpki-sever) or Prefix validation state, use the `clear bgp rpki` command.

Use the `server` option in the `clear bgp rpki` command in EXEC mode to clear the Resource Public Key Infrastructure (RPKI) server, purge Route Origin Authorizations (ROA) records downloaded from that server, renegotiate the TCP connection, and redownload ROA records.

Use `clear bgp rpki server WORD` command to clear a specific RPKI cache-server (hostname or IP address).

The `clear bgp rpki server *` command clears all RPKI cache-servers.

Use the `clear bgp rpki validation` command to re-run validation for all prefixes under the specified VRF or all VRFs, if not specified. If the address family is specified after the command, then only prefixes belonging to that address family will have their validation checks re-run.

clear bgp rpki server { * | WORD }

clear bgp { vrf <vrf name> } rpki validation { ipv4 | ipv6 }

Syntax Description

| | |
|-------------------|--|
| server | Clear RPKI cache-server. |
| validation | Clear RPKI origin-validation (trigger origin validation). |
| * | Clear all RPKI cache-servers. |
| WORD | Clear a specific RPKI cache-server (hostname or IP address). |
| ipv4 | Trigger origin validation walk for the IPv4 table. |
| ipv6 | Trigger origin validation walk for the IPv6 table. |

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

The following example shows how to configure BGP RPKI server:

```
router bgp 65000
  rpki server 10.0.101.2
    bind-source interface GigabitEthernet0/2/0/1
    transport tcp port 32001
    response-time 15
```

The following example shows the output of `show bgp rpki server summary` and `show bgp rpki table` commands:

```
RP/0/0/CPU0:router#show bgp rpki server summary
Wed May  6 10:26:42.532 PDT
```


| Hostname/Address | Transport | State | Time | ROAs (IPv4/IPv6) |
|------------------|-----------|-------|----------|------------------|
| 10.0.101.2 | TCP:32001 | ESTAB | 00:00:09 | 3/3 |

RP/0/0/CPU0:router#show bgp rpki table
Wed May 6 10:26:44.592 PDT

| Network | Maxlen | Origin-AS | Server |
|---------------|--------|-----------|------------|
| 1.1.1.0/24 | 32 | 9991 | 10.0.101.2 |
| 67.64.0.0/30 | 32 | 4002 | 10.0.101.2 |
| 67.128.0.0/30 | 32 | 4003 | 10.0.101.2 |

Processed 3 RPKI entries

The following example shows the output of show bgp rpki server summary command after clearing RPKI server 10.0.101.2:

RP/0/0/CPU0:router#clear bgp rpki server 10.0.101.2
Wed May 6 10:26:50.318 PDT

RP/0/0/CPU0:router#show bgp rpki server summary
Wed May 6 10:26:53.212 PDT

| Hostname/Address | Transport | State | Time | ROAs (IPv4/IPv6) |
|------------------|-----------|-------|----------|------------------|
| 10.0.101.2 | TCP:32001 | NONE | 00:00:04 | N/A |

RP/0/0/CPU0:router#show bgp rpki table
Wed May 6 10:26:55.837 PDT

Processed 0 RPKI entries

RP/0/0/CPU0:router#show bgp rpki server summary
Wed May 6 10:27:04.663 PDT

| Hostname/Address | Transport | State | Time | ROAs (IPv4/IPv6) |
|------------------|-----------|-------|----------|------------------|
| 10.0.101.2 | TCP:32001 | ESTAB | 00:00:11 | 0/0 |

RP/0/0/CPU0:router#show bgp rpki table
Wed May 6 10:27:17.513 PDT

| Network | Maxlen | Origin-AS | Server |
|---------------|--------|-----------|------------|
| 1.1.1.0/24 | 32 | 9991 | 10.0.101.2 |
| 67.64.0.0/30 | 32 | 4002 | 10.0.101.2 |
| 67.128.0.0/30 | 32 | 4003 | 10.0.101.2 |

Processed 3 RPKI entries

clear bgp self-originated

To clear Border Gateway Protocol (BGP) routes that are self-originated, use the **clear bgp self-originated** command in XR EXEC mode.

clear bgp self-originated

| | | |
|---------------------------|---|--|
| Syntax Description | ipv4 | Specifies IP Version 4 address prefixes. |
| | unicast | Specifies unicast address prefixes. |
| | labeled-unicast | Specifies labeled-unicast address prefixes. |
| | all | For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | all | For address family, specifies prefixes for all address families. |
| Command Default | No default behavior or values | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Self-originated routes are routes locally originated by the network command, redistribute command, or aggregate-address command. | |
| Task ID | Task ID | Operations |
| | bgp | execute |
| Examples | The following example shows how to clear self-originated IPv4 routes: | |
| | <pre>RP/0/RP0/CPU0:router# clear bgp ipv4 unicast self-originated</pre> | |
| Related Commands | Command | Description |
| | aggregate-address, on page 12 | Creates an aggregate entry in a BGP routing table. |

| Command | Description |
|---|---|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| redistribute (BGP), on page 169 | Redistributes routes from another routing protocol into BGP. |

clear bgp shutdown

To clear all Border Gateway Protocol (BGP) neighbors that shut down due to low memory, use the **clear bgp shutdown** command in XR EXEC mode.

clear bgp shutdown

| | |
|------------------------|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| labeled-unicast | Specifies labeled unicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |

Command Default No default behavior or values

Command Modes XR EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples The following example shows how to clear all shut-down BGP neighbors:

```
RP/0/RP0/CPU0:router# clear bgp shutdown
```

| Related Commands | Command | Description |
|------------------|---|--|
| | show bgp, on page 193 | Displays entries in the BGP routing table. |
| | show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |

clear bgp soft

To soft reset a group of Border Gateway Protocol (BGP) neighbors, use the **clear bgp soft** command in XR EXEC mode.

```
clear bgp { * ip-address | asas-number | external } soft [[ { in | { prefix-filter } | out } ]]
```

| Syntax Description | ipv4 | Specifies IP Version 4 address prefixes. |
|--------------------|---|---|
| | unicast | Specifies unicast address prefixes. |
| | labeled-unicast | Specifies labeled-unicast address prefixes. |
| | all | For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | Specifies IP Version 6 address prefixes. |
| | all | For address family, specifies prefixes for all address families. |
| | <i>nexthop-address</i> | Address of the next hop. |
| | * | Soft resets all BGP neighbors. |
| | <i>ip-address</i> | IP address of the neighbor to be reset. |
| | as <i>as-number</i> | Autonomous system (AS) number for all neighbors to be reset. Range for 2-byte numbers is 1 to 65535. Range for 4-byte numbers is 1.0 to 65535.65535. |
| | external | Specifies clearing of all external peers. |
| | in | (Optional) Triggers an inbound soft reset. If the in or out keyword is not specified, both inbound and outbound soft resets are triggered. |
| | prefix-filter | (Optional) Specifies to send a new Outbound Route Filter (ORF) to the neighbor. Neighbor installs the new ORF and resends its routes. |
| | out | (Optional) Triggers an outbound soft reset. If the in or out keyword is not specified, both inbound and outbound soft resets are triggered. |
| Command Default | No default behavior or value | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |

Use the **clear bgp soft** command to trigger a soft reset of the specified address families for the specified group of neighbors. This command is useful if you change the inbound or outbound policy for the neighbors, or any other configuration that affects the sending or receiving of routing updates.

If an outbound soft reset is triggered, BGP resends all routes for the address family to the given neighbors.

If an inbound soft reset is triggered, BGP by default sends a REFRESH request to the neighbor, if the neighbor has advertised the ROUTE_REFRESH capability. To determine whether the neighbor has advertised the ROUTE_REFRESH capability, use the **show bgp neighbors** command, and look for the following line of output:

```
Received route refresh capability from peer.
```

If the neighbor does not support route refresh, but the **soft-reconfiguration inbound** command is configured for the neighbor, then BGP uses the routes cached as a result of the **soft-reconfiguration inbound** command to perform the soft reset.

If you want BGP to use the cached routes even if the neighbor supports route refresh, you can use the **always** keyword when configuring the **soft-reconfiguration inbound** command.

If the neighbor does not support route refresh and the **soft-reconfiguration inbound** command is not configured, then inbound soft reset is not possible. In this case, an error is printed.



Note

By default, if the configuration for an inbound or outbound route policy is changed, BGP performs an automatic soft reset. Use the **bgp auto-policy-soft-reset disable** command to disable this behavior.

Task ID

Task ID Operations

bgp execute

Examples

The following example shows how to trigger an inbound soft clear for IPv4 unicast routes received from neighbor 10.0.0.1:

```
RP/0/RP0/CPU0:router# clear bgp ipv4 unicast 10.0.0.1 soft in
```

Related Commands

| Command | Description |
|--|---|
| bgp auto-policy-soft-reset disable, on page 29 | Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified. |
| clear bgp, on page 82 | Resets a group of BGP neighbors. |
| clear bgp self-originated, on page 98 | Clears self-originated routes. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

| Command | Description |
|---|--|
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |
| soft-reconfiguration inbound, on page 345 | Configures the software to store updates received from a neighbor. |

cluster-id

To configure the cluster for a neighbor, use the **cluster-id** command in an appropriate configuration mode. To remove the cluster, use the **no** form of this command.

cluster-id *cluster-id*
no cluster-id [*cluster-id*]

Syntax Description

cluster-id Cluster ID of the router acting as a route reflector; maximum of four bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295.

Command Default

A cluster ID is not configured.

Command Modes

Neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A single route reflector can support multiple clusters. A neighbor can be associated with one cluster only. And the corresponding cluster ID is configured in neighbor configuration mode. If the cluster ID is not configured for a neighbor and the neighbor is a route reflector client, then the neighbor is assigned to the default cluster.

A neighbor will be considered to be a route reflector client only if it is configured as a route reflector client in the appropriate address-family configuration mode.

Configuring the cluster ID using the **cluster-id** command for a neighbor group or session group under the neighbor group configuration mode or the session group configuration mode causes all neighbors using the group to inherit the characteristics configured with the command. Configuring the command directly for the neighbor overrides the value inherited from the group.

To increase redundancy and avoid a single point of failure in the network, the clusters might be connected to more than one route reflector. In this case, the neighbor to cluster-id mapping at all the route reflectors must be the same so that a route reflector can recognize updates from route reflectors that are connected to the same clusters.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure the local router as one of the route reflectors serving three clusters. Neighbor 192.168.70.25 is assigned to the default cluster with cluster ID 1.

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# bgp cluster-id 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# cluster-id 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client

RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.25
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client

RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.70.26
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RP0/CPU0:router(config-bgp-nbr)# cluster-id 3
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

Related Commands

| Command | Description |
|---|---|
| route-reflector-client, on page 180 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

cluster-id allow-equal

To enable a router to accept BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs, use the **cluster-id allow-equal** command in BGP neighbor address-family configuration mode. To revert to the default behaviour, use the **disable** option, or **no** form of this command.

cluster-id allow-equal [**disable**]

| Syntax Description | disable Specifies that BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs, is rejected. | | | | |
|---------------------------|--|---------|--------------|---------------|------------------------------|
| Command Default | A router rejects BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs. | | | | |
| Command Modes | BGP neighbor address-family configuration. | | | | |
| Command History | <table> <tr> <th>Release</th><th>Modification</th></tr> <tr> <td>Release 6.6.3</td><td>This command was introduced.</td></tr> </table> | Release | Modification | Release 6.6.3 | This command was introduced. |
| Release | Modification | | | | |
| Release 6.6.3 | This command was introduced. | | | | |
| Usage Guidelines | <p>To avoid routing loops, the cluster-id allow-equal configuration must be done with care.</p> <p>The following example shows how to enable a router to accept BGP routes which have the same first cluster-ID as the router's own cluster-ID in the list of cluster-IDs.</p> <pre>Router(config)# router bgp 100 Router(config-bgp)# neighbor 10.1.1.1 Router(config-bgp-nbr)# address-family ipv4 unicast Router(config-bgp-nbr-af)# cluster-id allow-equal</pre> <p>The following example shows how to disable a router from accepting BGP routes from the same cluster-ID as the router's own cluster-ID in the list of cluster-IDs.</p> <pre>Router(config)# router bgp 100 Router(config-bgp)# neighbor 10.1.1.1 Router(config-bgp-nbr)# address-family ipv4 unicast Router(config-bgp-nbr-af)# cluster-id allow-equal disable</pre> | | | | |

default-information originate (BGP)

To allow origination of a default route to be redistributed into the Border Gateway Protocol (BGP) from another protocol, use the **default-information originate** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

default-information originate
no default-information originate

| | |
|---------------------------|--|
| Syntax Description | This command has no arguments or keywords. |
|---------------------------|--|

| | |
|------------------------|---|
| Command Default | BGP does not permit redistribution of a default route into BGP. |
|------------------------|---|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|------------------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

Use the **redistribute** command to redistribute routes from another protocol into BGP. By default, if these routes include the default route (0.0.0.0/0 for IPv4 or ::/0 for IPv6), the default route is ignored. Use the **default-information originate** command to change this behavior so that the default route is not ignored and is redistributed into BGP along with the other routes for the protocol being redistributed.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to configure BGP to redistribute the default route into BGP:

```
RP/0/RP0/CPU0:router(config)#router bgp 164
RP/0/RP0/CPU0:router(config-bgp)# default-information originate
```


Related Commands

| Command | Description |
|---|--|
| redistribute (BGP), on page 169 | Redistributes routes from another protocol into BGP. |

default-metric (BGP)

To set default metric values for the Border Gateway Protocol (BGP), use the **default-metric** command in an appropriate configuration mode. To disable metric values, use the **no** form of this command.

default-metric *value*
no default-metric [*value*]

| | | |
|--|--|------------------------------|
| Syntax Description | <i>value</i> Default metric value appropriate for the specified routing protocol. Range is 1 to 4294967295. | |
| Command Default | A metric is not set. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the default-metric command to set the Multi Exit Discriminator (MED) to advertise to peers for routes that do not already have a metric set (routes that were received with no MED attribute). | |
|  Note | The metric values that you apply using the default-metric command take effect only for a new prefix which gets into the BGP table. The metrics for the existing prefixes in the BGP table remain the same. Also, when you remove the default-metric command from the configuration, the metrics which were previously assigned for prefixes are not updated. To get out of this condition, clear the BGP neighborship. | |
| | | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>The following example shows how to set the BGP default metric:</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 109 RP/0/RP0/CPU0:router(config-bgp)# default-metric 10</pre> | |

default-originate

To cause a Border Gateway Protocol (BGP) speaker (the local router) to send the default route 0.0.0.0/0 to a neighbor for use as a default route, use the **default-originate** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

default-originate [{**inheritance-disable** | **route-policy** *route-policy-name*}]
no default-originate [{**inheritance-disable** | **route-policy** *route-policy-name*}]

| Syntax Description | <div> inheritance-disable (Optional) Prevents the default-originate command characteristics from being inherited from a parent group. </div> <div> route-policy <i>route-policy-name</i> (Optional) Specifies the name of a route policy. The route policy allows route 0.0.0.0 to be injected conditionally. IPv6 address family is supported. </div> | | | | |
|---------------------------|---|---------|--------------|---------------|------------------------------|
| Command Default | The default route is not advertised to BGP neighbors. | | | | |
| Command Modes | IPv4 neighbor address family configuration IPv6 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration IPv4 address family group configuration IPv6 address family group configuration | | | | |
| Command History | <table> <tr> <th data-bbox="386 1167 943 1199">Release</th><th data-bbox="959 1167 1528 1199">Modification</th></tr> <tr> <td data-bbox="386 1230 943 1262">Release 5.0.0</td><td data-bbox="959 1230 1528 1262">This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>The default-originate command does not require the presence of the default route (0.0.0.0/0 for IPv4 or ::/0 for IPv6) in the local router. When the default-originate command is used with a route policy, the default route is advertised if any route in the BGP table matches the policy.</p> | | | | |
| Task ID | <table> <tr> <th data-bbox="386 1556 456 1587">Task ID</th><th data-bbox="472 1556 1528 1587">Operations</th></tr> <tr> <td data-bbox="386 1650 456 1682">bgp</td><td data-bbox="472 1650 1528 1713">read, write</td></tr> </table> | Task ID | Operations | bgp | read, write |
| Task ID | Operations | | | | |
| bgp | read, write | | | | |
| Examples | <p>The following example shows how to unconditionally advertise the route 0.0.0.0/0 to the neighbor 172.20.2.3:</p> | | | | |

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# default-originate
```

The following example shows how to advertise the route 0.0.0.0/0 to the neighbor 172.20.2.3 only if a route exists in the BGP table that matches the route policy called default-default-policy:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# default-originate route-policy default-default-policy
```

Related Commands

| Command | Description |
|--|---|
| default-information originate (BGP), on page 107 | Allows the default route to be redistributed into BGP from another routing protocol. |
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |

description (BGP)

To annotate a neighbor, neighbor group, or session group, use the **description** command in an appropriate configuration mode. To remove the annotation, use the **no** form of this command.

description *text*
no description [{*text*}]

Syntax Description

text Meaningful description or comment. Maximum of 80 characters.

Command Default

No comment or description exists.

Command Modes

Neighbor group configuration

Neighbor configuration

Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **description** command to provide a description of a neighbor, neighbor group, or session group. The description is used to save user comments and does not affect software function.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure the description “Our best customer” on the neighbor 192.168.13.4:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)#router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#neighbor 192.168.13.4
RP/0/RP0/CPU0:router(config-bgp-nbr)#description Our best customer
```

distance bgp

To allow the use of external, internal, and local administrative distances that could be used to prefer one class of routes over another, use the **distance bgp** command in an appropriate configuration mode. To disable the use of administrative distances, use the **nono** form of this command.

distance bgp *external-distance internal-distance local-distance*

no distance bgp [*external-distance internal-distance local-distance*]

Syntax Description

| | |
|--------------------------|--|
| <i>external-distance</i> | Administrative distance for Border Gateway Protocol (BGP) external routes. External routes are routes for which the best path is learned from a neighbor external to the autonomous system. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table. |
| <i>internal-distance</i> | Administrative distance for BGP internal routes. Internal routes are those routes that are learned from another BGP entity within the same autonomous system. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table. |
| <i>local-distance</i> | Administrative distance for BGP local routes. The <i>local-distance</i> argument applies to locally generated aggregate routes (such as the routes generated by the aggregate-address command) and backdoor routes installed in the routing table. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table. |

Command Default

external-distance : 20

internal-distance : 200

local-distance : 200

Command Modes

IPv4 address family configuration

IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **distance bgp** command if another protocol is known to be able to provide a better route to a node than was actually learned using external BGP, or if some internal routes should be preferred by BGP.



Note

Changing the administrative distance of BGP internal routes is considered risky and is not recommended. One problem that can arise is the accumulation of routing table inconsistencies, which can interfere with routing.

An administrative distance is a rating of the trustworthiness of a routing information source. Numerically, an administrative distance is an integer from 1 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored.

Task ID

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows that iBGP routes are preferable to locally generated routes, so the administrative distance values are set accordingly:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#distance bgp 20 20 200
```

Related Commands

| Command | Description |
|------------------|--|
| distance (IS-IS) | Defines the administrative distance assigned to routes discovered by the IS-IS protocol. |
| distance (OSPF) | Defines OSPF route administrative distances based on route type. |

dscp (BGP)

To set the differentiated services code point (DSCP) value, use the **dscp** command in the appropriate configuration mode. To remove the **dscp** command from the configuration file and restore the system to its default interval values, use the no form of this command.

dscp *value*
no dscp [{*value*}]

| Syntax Description | <i>value</i> Value of the DSCP. The DSCP value can be a number from 0 to 63, or it can be one of the following keywords: default , ef , af11 , af12 , af13 , af21 , af22 , af23 , af31 , af32 , af33 , af41 , af42 , af43 , cs1 , cs2 , cs3 , cs4 , cs5 , cs6 , or cs7 . | | | | |
|---------------------------|--|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | Neighbor configuration Neighbor session group configuration Neighbor group configuration | | | | |
| Command History | <table> <tr> <th>Release</th><th>Modification</th></tr> <tr> <td>Release 5.0.0</td><td>This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the dscp command to change the minimum and maximum packet thresholds for the DSCP value.</p> <p>Table 1: dscp Default Settings, on page 114 lists the DSCP default settings used by the dscp command. The DSCP value, corresponding minimum threshold, maximum threshold, and mark probability are listed. The last row of the table (the row labeled "default") shows the default settings used for any DSCP value not specifically shown in the table.</p> | | | | |

Table 1: dscp Default Settings

| DSCP (Precedence) | Minimum Threshold | Maximum Threshold | Mark Probability |
|-------------------|-------------------|-------------------|------------------|
| af11 | 32 | 40 | 1/10 |
| af12 | 28 | 40 | 1/10 |
| af13 | 24 | 40 | 1/10 |
| af21 | 32 | 40 | 1/10 |
| af22 | 28 | 40 | 1/10 |

| DSCP (Precedence) | Minimum Threshold | Maximum Threshold | Mark Probability |
|-------------------|-------------------|-------------------|------------------|
| af23 | 24 | 40 | 1/10 |
| af31 | 32 | 40 | 1/10 |
| af32 | 28 | 40 | 1/10 |
| af33 | 24 | 40 | 1/10 |
| af41 | 32 | 40 | 1/10 |
| af42 | 28 | 40 | 1/10 |
| af43 | 24 | 40 | 1/10 |
| cs1 | 22 | 40 | 1/10 |
| cs1 | 24 | 40 | 1/10 |
| cs3 | 26 | 40 | 1/10 |
| cs4 | 28 | 40 | 1/10 |
| cs5 | 30 | 40 | 1/10 |
| cs6 | 32 | 40 | 1/10 |
| cs7 | 34 | 40 | 1/10 |
| ef | 36 | 40 | 1/10 |
| default | 20 | 40 | 1/10 |

Task ID**Task ID Operations**

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to set the DSCP value to af32:

```
RP/0/RP0/CPU0:router(config)# router bgp 5
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# dscp af32
```

ebgp-multihop

To accept and attempt Border Gateway Protocol (BGP) connections to external peers residing on networks that are not directly connected, use the **ebgp-multihop** command in an appropriate configuration mode. To disable connections to external peers and allow only direct connections between neighbors, use the **no** form of this command.

ebgp-multihop [*{ttl-value}*] [**mpls**]
no ebgp-multihop [*{ttl-value}*] [**mpls**]

| Syntax Description | <div> <i>ttl-value</i> (Optional) Time-to-live (TTL) value. Range is 1 to 255 hops. </div> <div> mpls (Optional) Disables BGP label rewrite. </div> | | | | |
|---------------------------|---|---------|--------------|---------------|------------------------------|
| Command Default | Default TTL value is 255. | | | | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | | | | |
| Command History | <table> <tr> <th data-bbox="334 974 842 1014">Release</th><th data-bbox="842 974 1505 1014">Modification</th></tr> <tr> <td data-bbox="334 1031 842 1071">Release 5.0.0</td><td data-bbox="842 1031 1505 1071">This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the ebgp-multihop command to enable multihop peerings with external BGP neighbors. The BGP protocol states that external neighbors must be directly connected (one hop away). The software enforces this by default; however, the ebgp-multihop command can be used to override this behavior.</p> <p>Use of the mpls option in the ebgp-multihop command prevents BGP from enabling MPLS on the peering interface and also prevents allocation of Implicit-NULL rewrite labels for nexthop addresses learned from the peer. This is useful in some scenarios in which MPLS forwarding labels to the nexthops have already been learned via BGP labeled-unicast or LDP.</p> <p>If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.</p> | | | | |
| Task ID | <table> <tr> <th data-bbox="334 1577 423 1617">Task ID</th><th data-bbox="423 1577 1505 1617">Operations</th></tr> <tr> <td data-bbox="334 1671 423 1711">bgp</td><td data-bbox="423 1671 1505 1711">read, write</td></tr> </table> | Task ID | Operations | bgp | read, write |
| Task ID | Operations | | | | |
| bgp | read, write | | | | |

Examples

The following example shows how to allow a BGP connection to neighbor 172.20.16.6 of up to 255 hops away:

```
RP/0/RP0/CPU0:router(config)# router bgp 109  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6  
RP/0/RP0/CPU0:router(config-bgp-nbr)# ebgp-multihop
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

enforce-first-as

To enable the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, use the **enforce-first-as** command in an appropriate configuration mode. To disable enforcing the first autonomous system in the AS path of a route received from an eBGP peer to be the same as the remote autonomous system, use the **no** form of this command.

enforce-first-as
no enforce-first-as

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | This command has no arguments or keywords. | |
| Command Default | By default, the software requires the first autonomous system (in the AS path) of a route received from an eBGP peer to be the same as the remote autonomous system configured. | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>By default, the software ignores any update received from an eBGP neighbor that does not have the autonomous system configured for that neighbor at the beginning of the AS path. When configured, the command applies to all eBGP peers under the neighbor, neighbor group or session group.</p> <p>At any given time, either the enforce-first-as command or the enforce-first-as-disable, on page 120 command can be configured under a given neighbor, neighbor group or session group. Configuring one command overwrites the other command.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>The following example shows a configuration in which incoming updates from eBGP neighbors are checked to ensure the first AS number in the AS path is the same as the configured AS number for the neighbor:</p> <pre>RP/0/RP0/CPU0:router(config)# router bgp 100</pre> | |

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4  
RP/0/RP0/CPU0:router(config-bgp-nbr)# enforce-first-as
```

Related Commands

| Command | Description |
|--|--|
| bgp enforce-first-as disable, on page 55 | Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in router configuration mode. |
| enforce-first-as-disable, on page 120 | Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |

enforce-first-as-disable

To disable the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, use the **enforce-first-as-disable** command in an appropriate configuration mode. To re-enable enforcing first autonomous system in the AS path of a route received from an eBGP peer to be the same as the remote autonomous system, use the **no** form of this command.

enforce-first-as-disable

no enforce-first-as-disable

Syntax Description

This command has no arguments or keywords.

Command Default

By default, the software requires the first autonomous system (in the AS path) of a route received from an eBGP peer to be the same as the remote autonomous system configured.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, the software ignores any update received from an eBGP neighbor that does not have the autonomous system configured for that neighbor at the beginning of the AS path. When configured, the command applies to all eBGP peers under the neighbor, neighbor-group or session-group.

At any given time, either the **enforce-first-as-disable** command or the [enforce-first-as, on page 118](#) command can be configured under a given neighbor, neighbor group or session group. Configuring one command overwrites the other command.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows a configuration in which incoming updates from eBGP neighbors are not checked to ensure the first AS number in the AS path is the same as the configured AS number for the neighbor:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
```



```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4  
RP/0/RP0/CPU0:router(config-bgp-nbr)# enforce-first-as-disable
```

Related Commands

| Command | Description |
|--|--|
| bgp enforce-first-as disable, on page 55 | Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in router configuration mode. |
| enforce-first-as, on page 118 | Enables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, under neighbor configuration mode, neighbor group configuration mode, and session group configuration mode. |

ibgp policy out enforce-modifications

To allow an outbound route policy for an internal BGP (iBGP) peer to modify all BGP route attributes, only when an iBGP route is sent to another iBGP peer (only on route-reflectors), use the **ibgp policy out enforce-modifications** command in router configuration mode. To disable this feature, use the **no** form of this command.

ibgp policy out enforce-modifications
no ibgp policy out enforce-modifications

Syntax Description This command has no arguments or keywords.

Command Default ibgp policy out enforce-modifications is disabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **ibgp policy out enforce-modifications** command to set and modify BGP route attributes for updates to iBGP peers.

If the **ibgp policy out enforce-modifications** command is configured under router BGP configuration, then all the changes made by the outbound policy for an iBGP peer will be present in an update message sent to the peer.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to set the ibgp policy out enforce-modifications:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 6500
RP/0/RP0/CPU0:router(config-bgp)# ibgp policy out enforce-modifications
```

ignore-connected-check

To enable the software to bypass the directly connected next hop check for single-hop eBGP peering, use the **ignore-connected-check** command in an appropriate configuration mode. To re-enable the directly connected next hop check, use the **no** form of this command.


ignore-connected-check [{**inheritance-disable**}]
no ignore-connected-check

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | inheritance-disable Prevents the ignore-connected-check command from being inherited from the parent. | |
| Command Default | Ability to bypass the directly connected next hop check is disabled. | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | The following example shows how to enable ignore-connected check configuration for neighbor 10.2.3.4: RP/0/RP0/CPU0:router(config)# router bgp 100 RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.2.3.4 RP/0/RP0/CPU0:router(config-bgp-nbr)# ignore-connected-check | |

keychain

To apply key chain-based authentication on a TCP connection between two Border Gateway Protocol (BGP) neighbors, use the **keychain** command in an appropriate configuration mode. To disable key chain authentication, use the **no** form of this command.

keychain *name*
no keychain [{*name*}]

| Syntax Description | <i>name</i> Key chain name configured using the keychain command. The name must be a maximum of 32 alphanumeric characters. | | | | |
|---|---|---------|--------------|---------------|------------------------------|
| Command Default | When this command is not specified in the appropriate configuration mode, key chain authentication is not enabled on a TCP connection between two BGP neighbors. | | | | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | | | | |
| Command History | <table> <tr> <th>Release</th><th>Modification</th></tr> <tr> <td>Release 5.0.0</td><td>This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Specify a key chain to enable key chain authentication between two BGP peers. Use the keychain command to implement hitless key rollover for authentication.</p> <p>If this command is configured for a neighbor group or a session group, a neighbor using the group inherits the configuration. Values of commands configured specifically for a neighbor override inherited values.</p> | | | | |
|  | | | | | |
| Note | BGP only supports HMAC-MD5 and HMAC-SHA1-12 cryptographic algorithms. | | | | |
| Task ID | <table> <tr> <th>Task ID</th><th>Operations</th></tr> <tr> <td>bgp</td><td>read, write</td></tr> </table> | Task ID | Operations | bgp | read, write |
| Task ID | Operations | | | | |
| bgp | read, write | | | | |
| Examples | The following example shows how to configure neighbor 172.20.1.1 to use the key chain authentication configured in the keychain_A key chain: | | | | |

```
RP/0/RP0/CPU0:router(config)# router bgp 140  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain keychain_A
```

Related Commands

| Command | Description |
|---|---|
| keychain-disable, on page 126 | Overrides any inherited key chain configuration from a neighbor group or session group for BGP neighbors. |

keychain-disable

**Note**

Effective with Release 3.9.0, the **keychain-disable** command was replaced by the **keychain inheritance-disable** command. See the [keychain inheritance-disable, on page 128](#) command for more information.

To override any inherited key chain configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **keychain-disable** command in an appropriate configuration mode. To disable overriding any inherited key chain command, use the **no** form of this command.

keychain-disable
no keychain-disable

Syntax Description

This command has no arguments or keywords.

Command Default

Configured key chains for neighbor and session groups are inherited.

Command Modes

Neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different **keychain** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **keychain-disable** on a neighbor that uses the group disables key chain authentication for the neighbor.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
```

```
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)#use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain-disable
```

Related Commands

| Command | Description |
|---------------------------------------|---|
| keychain, on page 124 | Enables key chain authentication on a TCP connection between two BGP neighbors. |

keychain inheritance-disable

To override any inherited key chain configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **keychain inheritance-disable** command in an appropriate configuration mode. To disable overriding any inherited key chain command, use the **no** form of this command.

keychain inheritance-disable
no keychain inheritance-disable

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | This command has no arguments or keywords. | |
| Command Default | Configured key chains for neighbor and session groups are inherited. | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>If you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different keychain command specifically on a neighbor that uses the group overrides the inherited value. Specifying keychain inheritance-disable on a neighbor that uses the group disables key chain authentication for the neighbor.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | <p>The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:</p> <pre>RP/0/RP0/CPU0:router(config)#router bgp 140 RP/0/RP0/CPU0:router(config-bgp)# session-group group1 RP/0/RP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1 RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2 RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1 RP/0/RP0/CPU0:router(config-bgp-nbr)# keychain inheritance-disable</pre> | |

Related Commands

| Command | Description |
|---------------------------------------|---|
| keychain, on page 124 | Enables key chain authentication on a TCP connection between two BGP neighbors. |

local-as

To allow customization of the autonomous system number for external Border Gateway Protocol (eBGP) neighbor peerings, use the **local-as** command in an appropriate configuration mode. To disable customization of local autonomous system values for eBGP neighbor peerings, use the **no** form of this command.

local-as {*as-number* [**no-prepend** [**replace-as** [**dual-as**]]] | **inheritance-disable**}
no local-as [{*as-number* [**no-prepend** [**replace-as** [**dual-as**]]] | **inheritance-disable**}]

| | | |
|---------------------------|--|---|
| Syntax Description | <i>as-number</i> | Valid autonomous system number. Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. Cannot be the autonomous system number to which the neighbor belongs. |
| | no-prepend | (Optional) Specifies that local autonomous system values are not prepended to announcements from the neighbor. |
| | replace-as | (Optional) Specifies that prepend only local autonomous system values to announcements to the neighbor. |
| | dual-as | (Optional) Dual-AS mode. |
| | inheritance-disable | Prevents local AS from being inherited from the parent. |
| Command Default | The BGP autonomous system number specified in the router bgp command is used, except when confederations are in use. The confederation autonomous system is used for external neighbors in an autonomous system that is not part of the confederation. | |
| Command Modes | Neighbor configuration | |
| | Neighbor group configuration | |
| | Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>You can specify the autonomous system number the local BGP uses to peer with each neighbor. The autonomous system number specified with this command cannot be the local BGP autonomous system number (specified</p> | |

with the **router bgp** command) or the autonomous system number of the neighbor (specified with the **remote-as** command). This command cannot be specified for internal neighbors or for external neighbors in an autonomous system that is part of a confederation.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

In case the router is an ASBR and adding the **local-as** configuration changes the BGP neighbor from being an eBGP to iBGP, incremental addition of the **local-as** command is not supported. Remove the neighbor configuration, then configure the entire neighbor, including **remote-as** configuration and the new **local-as** configuration, in one commit.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows BGP using autonomous system 30 for the purpose of peering with neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 300
RP/0/RP0/CPU0:router(config-bgp-nbr)# local-as 30
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

maximum-paths (BGP)

To control the maximum number of parallel routes that Border Gateway Protocol (BGP) installs in the routing table, use the **maximum-paths** command in an appropriate configuration mode. To set the maximum number of parallel routes the software installs to the default value, use the **no** form of this command.

```
maximum-paths {ebgp | ibgp | eibgp} maximum [{unequal-cost}] [{selective}]
no maximum-paths {ebgp | ibgp | eibgp} [{maximum}] [{unequal-cost}]
```

| | | |
|--------------------|---------------------|---|
| Syntax Description | ebgp | Specifies external BGP multipath peers. |
| | ibgp | Specifies internal BGP multipath peers. |
| | eibgp | Specifies internal and external BGP multipath peers. eiBGP allows simultaneous use of internal and external paths. |
| | <i>maximum</i> | Maximum number of parallel routes that BGP installs in the routing table. Range is 2 to 32. |
| | unequal-cost | (Optional) Allows iBGP multipaths to have different BGP next-hop Interior Gateway Protocol (IGP) metrics. This option is available when the ibgp keyword is used. |
| | selective | (Optional) Allows BGP to be configured such that only routes from selected neighbors can be considered for multipath. This option is used with the multipath option. |

Command Default One path is installed in the routing table.

Command Modes IPv4 address family configuration
IPv6 address family configuration

| | | |
|-----------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-paths** command to allow the BGP to allow the BGP protocol to install multiple paths into the routing table for each prefix. With the eBGP option, multiple paths are installed for external peers that are from the same autonomous system and are equal cost (according to the BGP best-path algorithm). Similarly with the iBGP option, multiple paths are installed for internal peers that are equal cost based on the BGP best-path algorithm. With the eiBGP option, multiple paths from both iBGP and eBGP are eligible for multipath selection. The IGP metric to the BGP next hop is the same as the best-path IGP metric unless the router is configured for unequal cost iBGP multipath or eiBGP multipath. The **selective** option restricts multipath eligible routes to those that come from peers configured with the **multipath** option.

See *Implementing BGP* in the *Routing Configuration Guide for Cisco NCS 6000 Series Routers* for information on the BGP best-path algorithm.



Note The **maximum-paths** command with the **eibgp** keyword cannot be configured if the **ibgp** or **ebgp** keywords have been configured, because the **eibgp** keyword is a super set of the **ibgp** or **ebgp** keywords.

Task ID

Task ID Operations

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to allow a maximum of four paths to a destination to be installed into the IPv4 unicast routing table:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 4
RP/0/RP0/CPU0:router(config-bgp-af)# commit
```

The following example shows how you can configure selective multipath for iBGP and eBGP peers.



Note This configuration requires the **multipath** option to be configured for the neighbors. See the **multipath** command in the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference Guide* for more information.

For information on how this configuration is used, see the BGP Selective Multipath section in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ibgp 4 selective
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 5 selective
RP/0/RP0/CPU0:router(config-bgp-af)# commit
```

maximum-prefix (BGP)

To control how many prefixes can be received from a neighbor, use the **maximum-prefix** command in an appropriate configuration mode. To set the prefix limits to the default values, use the **no** form of this command.

maximum-prefix *maximum* [{*threshold*}] [{**warning-only**}] [**restart** *time-interval*]

no maximum-prefix *maximum* [{*threshold*}] [{**warning-only**}] [**restart** *time-interval*]

Syntax Description

| | |
|-------------------------------------|---|
| <i>maximum</i> | Maximum number of prefixes allowed from this neighbor. Range is from 1 to 4294967295. |
| | Note When using additional-paths feature, each path with a unique path ID received from a peer is counted separately for the purpose of maximum-prefix functionality. Hence, the <i>maximum</i> value should be configured appropriately when the peer is capable of sending additional-paths. |
| <i>threshold</i> | (Optional) Integer specifying at what percentage of the <i>maximum</i> argument value the software starts to generate a warning message. Range is from 1 to 100. |
| warning-only | (Optional) Instructs the software to only generate a log message when the <i>maximum</i> argument value is exceeded, and not to terminate the peering. |
| restart <i>time-interval</i> | (Optional) Sets the time interval (in minutes) after which peering session should be reestablished. Configure restart time interval in minutes. Range is from 1 to 65535. |

Command Default

When this command is not specified, the following defaults apply:

- IPv4 Unicast: 1048576

- IPv4 Labeled-unicast: 131072
- IPv6 Unicast: 524288
- IPv6 Labeled-unicast: 131072
- IPv4 Tunnel: 1048576
- IPv4 Multicast: 131072
- IPv6 Multicast: 131072
- IPv4 MVPN: 2097152
- VPNv4 Unicast: 2097152
- IPv4 MDT: 131072
- VPNv6 Unicast: 1048576
- L2VPN EVPN: 2097152
- IPv4 Flowspec: 1048576
- IPv6 Flowspec: 524288
- VPNv4 Flowspec: 2097152
- VPNv6 Flowspec: 1048576

The default threshold, when a warning message is generated, is 75 percent.

Command Modes

IPv4 address family group, neighbor address family, and neighbor group address family configuration

IPv6 address family group, neighbor address family, and neighbor group address family configuration

IPv4 tunnel address family group, neighbor group address family, and neighbor address family configuration

IPv4 flowspec under neighbor address family, neighbor group address family, and address family group configuration

IPv6 flowspec under neighbor address family, neighbor group address family, and address family group configuration

VPNv4 flowspec under neighbor address family, neighbor group address family, and address family group configuration

VPNv6 flowspec under neighbor address family, neighbor group address family, and address family group configuration

L2VPN EVPN under neighbor address family, neighbor group address family, and address family group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.2.1 | This command was introduced. |

Usage Guidelines

Use the **maximum-prefix** command to configure a maximum number of prefixes that a BGP router is allowed to receive from a neighbor. It adds another mechanism (besides routing policy) to control prefixes received from a peer.

When the number of received prefixes exceeds the maximum number configured, the software terminates the peering, by default, after sending a cease notification to the neighbor. However, if the **warning-only** keyword is configured, the software writes only a log message, but continues peering with the sender. If the peer is terminated, the peer stays down until the **clear bgp** command is issued or the **restart time-interval** option is used.

This command takes effect immediately if configured on an established neighbor, unless the number of prefixes received from the neighbor already exceeds the configured limits.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

This example shows the maximum number of IP Version 6 (IPv6) unicast prefixes allowed from neighbor 192.168.40.25 set to 5000, threshold value 80%, and restart time interval 20 minutes:

```
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 192.168.40.25
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#maximum-prefix 5000 80 restart 20
```

This example shows the maximum number of IP Version 4 (IPv4) unicast prefixes allowed from the neighbor 192.168.40.24 set to 1000:

```
RP/0/RP0/CPU0:router(config-bgp)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# maximum-prefix 1000
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| clear bgp, on page 82 | Resets a BGP connection using BGP hard or soft reconfiguration. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |

mpls activate (BGP)

To enable Multiprotocol Label Switching (MPLS) on an interface basis for ASBR and CSC configurations whenever a bgp confederation configuration is used, use the **mpls activate** command in bgp configuration mode. This is needed for InterAS (option B and C) and Carrier Supporting Carrier (CSC) configurations with confederations.

The normal InterAS and CSC configurations (without confederations) do not need to enable this.

To restore the system to its default condition, use the **no** form of this command.

mpls activate *interface id*
no mpls activate *interface id*

| | | |
|--------------------|--|--|
| Syntax Description | <i>interface id</i> Name of the interface. | |
| Command Default | No default behavior or values | |
| Command Modes | Router configuration Neighbor configuration IPv4 address family group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Thempls activate command enables MPLS on the interface specified and also adds the implicit null rewrite corresponding to the peer associated with the interface. The interface specified must be the one corresponding to the inter-AS ASBR or CSC peer.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Related Commands | Command | Description |
| | address-family (BGP), on page 6 | Enters address family configuration mode for configuring BGP routing sessions. |

multipath

Enables multiple paths for a BGP neighbor.

To disable this function, use the **no** form of this command.

multipath
no multipath

Command Default Multipath is disabled by default.

Command Modes Router BGP neighbor configuration

| Command History | Release | Modification |
|-----------------|-------------|------------------------------|
| | Release 4.2 | This command was introduced. |

Usage Guidelines To configure BGP selective multipath feature, the **multipath** option must be enabled on the required BGP neighbor. The **multipath** configuration for a neighbor works when configured with the **selective** option of the **maximum-paths** command.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | BGP | read, write |



Examples The following example shows how to enable multiple paths for a BGP neighbor.

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ibgp 4 selective
RP/0/RP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 5 selective
RP/0/RP0/CPU0:router(config-bgp-af)# neighbor 1.1.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# multipath
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# commit
```

neighbor (BGP)

To enter neighbor configuration mode for configuring Border Gateway Protocol (BGP) routing sessions, use the **neighbor** command in an appropriate configuration mode. To delete all configuration for a neighbor and terminate peering sessions with the neighbor, use the **no** form of this command.

neighbor *ip-address*
no neighbor *ip-address*

| | | |
|--|--|------------------------------|
| Syntax Description | <i>ip-address</i> IPv4 or IPv6 IP address of the BGP-speaking neighbor. | |
| Command Default | Neighbor mode is not specified. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | From router configuration mode, you can use this command to enter neighbor configuration mode. | |
| | From neighbor configuration mode, you can enter address family configuration for the neighbor by using the address-family command, which allows you to configure routing sessions for IP Version 4 and IP Version 6 address prefixes. | |
| | The neighbor command does not cause the neighbor to be configured and does not result in a peering to be established with the neighbor. To create the neighbor, you configure a remote autonomous system number by entering the remote-as command, or the neighbor can inherit a remote autonomous system from a neighbor group or session group if the use command is applied. | |
|  Note | A neighbor must have must a remote autonomous system number, and an IP address and address family must be enabled on the neighbor. | |
| | Unlike IPv4, IPv6 must be enabled before any IPv6 neighbors can be defined. Enable IPv6 in router configuration mode using the address-family command. | |
|  Note | Configuration for the neighbor cannot occur (peering is not established) until the neighbor is given a remote as-number and neighbor address. | |
| | The no form of this command causes the peering with the neighbor to be terminated and all configuration that relates to the neighbor to be removed. | |

Task ID

Task ID Operations

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to place the router in neighbor configuration mode for BGP routing process 1 and configure the neighbor IP address 172.168.40.24 as a BGP peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65000
```

The following example shows how to enable IPv6 for BGP, then place the router in neighbor configuration mode for an IPv6 neighbor, 3000::1, and configure neighbor 3000::1 as a BGP peer:

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv6 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 3000::1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2002
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv6 unicast
```

Related Commands

| Command | Description |
|---|---|
| address-family (BGP), on page 6 | Enters address family configuration mode for configuring BGP routing sessions. |
| remote-as (BGP), on page 173 | Adds an entry to the BGP neighbor table. |
| use, on page 368 | Inherits characteristics from a neighbor group, session group, or address family group. |

neighbor-group

To create a neighbor group and enter neighbor group configuration mode, use the **neighbor-group** command in router configuration mode. To remove a neighbor group and delete all configuration associated with the group, use the **no** form of this command.

neighbor-group *name*
no neighbor-group *name*

Syntax Description

name Neighbor group name.

Command Default

No neighbor group mode is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **neighbor-group** command puts the router in neighbor group configuration mode and creates a neighbor group.

A neighbor group helps you apply the same configuration to one or more neighbors. After a neighbor group is configured, each neighbor can inherit the configuration through the **use** command. If a neighbor is configured to use a neighbor group, the neighbor, by default, inherits the entire configuration of the neighbor group, which includes the address family-independent and address family-specific configurations. The inherited configuration can be overridden if you directly configure commands for the neighbor or if you configure session groups or address family groups with the **use** command.

From neighbor group configuration mode, you can configure address family-independent parameters for the neighbor group. To enter address family-specific configuration for the neighbor group, use the **address-family** command when in the neighbor group configuration mode.



Note

If an address family is configured for a neighbor group, neighbors that use the neighbor group attempt to exchange routes in that address family.

The **no** form of this command ordinarily causes all configuration for the neighbor group to be removed. If using the **no** form would result in a neighbor losing its remote autonomous system number, the configuration is rejected. In this scenario, the neighbor configuration must be either removed or configured with a remote autonomous system number before the neighbor group configuration can be removed.

**Note**

Neighbor groups should not be configured with a mixture of IPv4 and IPv6 address families, because such a neighbor group is not usable by any neighbor. Note that within the Cisco IOS XR system configuration architecture, it is possible to create such a neighbor group; however, any attempt to use it is rejected.

Task ID**Task ID** **Operations**

bgp read,
write

Related Commands

| Command | Description |
|---|--|
| address-family (BGP), on page 6 | Enters various address family configuration modes for configuring BGP routing sessions. |
| neighbor (BGP), on page 139 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| use, on page 368 | Inherits characteristics from a neighbor group, a session group, or an address family group. |

network (BGP)

To specify that the Border Gateway Protocol (BGP) routing process should originate and advertise a locally known network to its neighbors, use the **network** command in an appropriate configuration mode. To disable originating or advertising the network to neighbors, use the **no** form of this command.

network {*ip-address/prefix-length ip-address mask*} [**route-policy** *route-policy-name*]
no network {*ip-address/prefix-length ip-address mask*} [**route-policy** *route-policy-name*]

| | | |
|---------------------------|--|--|
| Syntax Description | <i>ip- address</i> | Network that BGP advertises. |
| | <i>/ prefix-length</i> | Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| | <i>ip-address mask</i> | Network mask applied to the <i>ip-address</i> argument. |
| | route-policy <i>route-policy-name</i> | (Optional) Specifies a route policy to use to modify the attributes of the network. |
| Command Default | No networks are specified. | |
| Command Modes | IPv4 address family configuration | |
| | IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | A network specified with this command is originated and advertised to neighbors only if there exists a route for the network in the routing table. That is, there must be a route learned using local or connected networks, static routing, or a dynamic IGP such as IS-IS or OSPF. | |
| | Other than the available system resources on the router, no limit exists on the number of network commands that can be configured. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to configure the local router to originate the IPv4 unicast network 172.20.0.0/16:

```
RP/0/RP0/CPU0:router(config)#router bgp 120  
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-af)# network 172.20.0.0/16
```

Related Commands

| Command | Description |
|---|---|
| network backdoor, on page 145 | Specifies a backdoor route to a BGP border router that provides better information about the network. |
| redistribute (BGP), on page 169 | Redistributes routes from one routing domain into another routing domain. |

network backdoor

To set the administrative distance on an external Border Gateway Protocol (eBGP) route to that of a locally sourced BGP route, causing it to be less preferred than an Interior Gateway Protocol (IGP) route, use the **network backdoor** command in an appropriate configuration mode. To disable setting the administrative distance to the value for locally sourced BGP routes, use the **no** form of this command.

network {*ip-address/prefix-length ip-address mask*} **backdoor**
no network {*ip-address/prefix-length ip-address mask*} **backdoor**

| | | |
|--------------------|---|--|
| Syntax Description | <i>ip-address</i> | Network that provides a backdoor route. |
| | <i>/prefix-length</i> | Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| | <i>mask</i> | Network mask applied to the <i>ip-address</i> argument. |
| Command Default | No backdoor routes are installed. | |
| Command Modes | IPv4 address family configuration | |
| | IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Configuring the network backdoor command does not cause BGP to originate a network, even if an IGP route for the network exists. Ordinarily, the backdoor network would be learned through both an eBGP and IGP. The BGP best-path selection algorithm does not change when a network is configured as a backdoor network. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | The following example shows IP Version 4 (IPv4) unicast network 192.168.40.0/24 configured as a backdoor network: RP/0/RP0/CPU0:router(config)# router bgp 109 RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast | |

```
RP/0/RP0/CPU0:router(config-bgp-af)# network 192.168.40.0/24 backdoor
```

Related Commands

| Command | Description |
|--|---|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |

next-hop-self

To disable next-hop calculation and insert your own address in the next-hop field of Border Gateway Protocol (BGP) updates, use the **next-hop-self** command in an appropriate configuration mode. To enable next-hop calculation, use the **no** form of this command.

next-hop-self [{**inheritance-disable**}]
no next-hop-self [{**inheritance-disable**}]

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | inheritance-disable (Optional) Allows a next-hop calculation override when this feature may be inherited from a neighbor group or address family group. | |
| Command Default | When this command is not specified, the software calculates the next hop for BGP updates accepted by the router. | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration IPv4 labeled-unicast address family configuration IPv6 labeled-unicast address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the next-hop-self command to set the BGP next-hop attribute of routes being advertised over a peering session to the local source address of the session.</p> <p>This command is useful in nonmeshed networks in which BGP neighbors may not have direct access to all other neighbors on the same IP subnet.</p> <p>If this command is configured for a neighbor group or address family group, a neighbor using the group inherits the configuration. Configuring the command specifically for a neighbor overrides any inherited value.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to set the next hop of the update field for all IP Version 4 (IPv4) unicast routes advertised to neighbor 172.20.1.1 to an address of the local router:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self
```

The following example shows how to disable the **next-hop-self** command for neighbor 172.20.1.1. If not overridden, the next hop would be inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-self
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# next-hop-self inheritance-disable
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| use, on page 368 | Inherits characteristics from a neighbor group, session group, or address family group. |

next-hop-unchanged

To disable overwriting of the next hop before advertising to external Border Gateway Protocol (eBGP) peers, use the **next-hop-unchanged** command in an appropriate configuration mode. To enable overwriting of the next hop, use the **no** form of this command.

next-hop-unchanged [{inheritance-disable}]
no next-hop-unchanged [{inheritance-disable}]

| | |
|---------------------------|---|
| Syntax Description | inheritance-disable (Optional) Allows overwriting of the next hop before advertising to eBGP peers when this feature may be inherited from a neighbor group or address family group. |
|---------------------------|---|

| | |
|------------------------|---|
| Command Default | Overwriting of the next hop is allowed. |
|------------------------|---|

| | |
|----------------------|--|
| Command Modes | IPv4 labeled-unicast address family configuration IPv6 labeled-unicast address family configuration IPv4 address family configuration IPv6 address family configuration |
|----------------------|--|

| Command History | <table border="1"><tr><th>Release</th><th>Modification</th></tr><tr><td>Release 5.0.0</td><td>This command was introduced.</td></tr></table> | Release | Modification | Release 5.0.0 | This command was introduced. |
|------------------------|--|---------|--------------|---------------|------------------------------|
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |

| | |
|-------------------------|---|
| Usage Guidelines | Use the next-hop-unchanged command to propagate the next hop unchanged for multihop eBGP peering sessions. This command should not be configured on a route reflector, and the next-hop-self command should not be used to modify the next-hop attribute for a route reflector when this feature is enabled for a route reflector client. |
|-------------------------|---|



| | |
|-------------|--|
| Note | Incorrectly setting BGP attributes for a route reflector can cause inconsistent routing, routing loops, or a loss of connectivity. Setting BGP attributes for a route reflector should be attempted only by an experienced network operator. |
|-------------|--|

| Task ID | <table border="1"><tr><th>Task ID</th><th>Operations</th></tr><tr><td>bgp</td><td>read, write</td></tr></table> | Task ID | Operations | bgp | read, write |
|----------------|---|---------|------------|-----|-------------|
| Task ID | Operations | | | | |
| bgp | read, write | | | | |

Examples

The following example shows how to disable the overwriting of next hops before advertising to eBGP peers:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
```

```
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-afgrp)# next-hop-unchanged disable  
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
```

Related Commands

| Command | Description |
|--|---|
| next-hop-self, on page 147 | Disables next-hop calculation and allows you to insert your own address in the next-hop field of BGP updates. |
| use, on page 368 | Inherits characteristics from a neighbor group, session group, or address family group. |

nexthop route-policy

To specify that BGP routes are resolved using only next hops whose routes match specific characteristics, use the **nexthop route-policy** command in the appropriate configuration mode. To remove the **nexthop route-policy** command from the configuration file and restore the system to its default behavior, use the **no** form of this command.

nexthop route-policy *route-policy-name*
no nexthop route-policy *route-policy-name*

Syntax Description

route-policy-name Route policy to use for filtering based on next hops.

Command Default

No default behavior or values

Command Modes

IPv4 address family configuration
IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

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

The BGP next-hop tracking feature allows you to specify that BGP routes are resolved using only next hops whose routes have the following characteristics:

- To avoid the aggregate routes, the prefix length must be greater than a specified value.
- The source protocol must be from a selected list, ensuring that BGP routes are not used to resolve next hops that could lead to oscillation.

This route policy filtering is possible because RIB identifies the source protocol of a route that resolves a next hop as well as the mask length associated with the route.

The next-hop attach point supports matching using the protocol name and mask length. BGP marks all next hops that are rejected by the route policy as invalid, and no best path is calculated for the routes that use the invalid next hop. The invalid next hops continue to stay in the active cache and can be displayed as part of the **show bgp nexthop** command with an invalid status.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to specify the route policy nexthop_A as the policy to use for filtering next hops:

```
RP/0/RP0/CPU0:router(config)# router bgp 109  
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-af)# nexthop route-policy nexthop_A
```


Related Commands

| Command | Description |
|--|--|
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |
| show bgp nexthops, on page 275 | Display statistical information about the BGP next hops. |

nexthop trigger-delay

To specify the delay for triggering next-hop calculations, use the **nexthop trigger-delay** command in the appropriate configuration mode. 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 | Specifies critical next-hop events. For example, when the next hop is unreachable. |
| | <i>delay</i> | Trigger delay, in milliseconds. Range is 0 to 4294967295. |
| | non-critical | Specifies noncritical next-hop events. For example, Interior Gateway Protocol (IGP) metric changes. |
| Command Default | critical : 3000 msec for IPv4 address family and IPv6 address family | |
| | critical : 0 msec for VPNv4 address family and VPNv6 address family | |
| | non-critical : 10000 msec IPv4, IPv6, VPNv4, and VPNv6 address families | |
| Command Modes | IPv4 address family configuration | |
| | Pv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the nexthop trigger-delay command to allow for a dynamic way for Interior Gateway Protocol (IGP) to converge. This convergence allows BGP to accumulate all notifications and trigger fewer walks, resulting in fewer interprocess communications (IPCs) to the Routing Information Base (RIB) for route addition, deletion, and modification and fewer updates to peers. | |
|  Note | A high <i>delay</i> value can be configured to effectively turn off next-hop tracking. | |
| | The non-critical <i>delay</i> value must always be set to at least equal or greater than the critical <i>delay</i> value. | |
| | The <i>delay</i> should be slightly higher than the time it takes for the IGP to settle into a steady state after some event (IGP convergence time). | |

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to set the critical next-hop trigger delay to 3500 milliseconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 109  
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-af)# nexthop trigger-delay critical 3500
```

nsr (BGP)

To activate Border Gateway Protocol (BGP) nonstop routing (NSR), use the **nsr** command in BGP XR Config mode. To deactivate BGP NSR, use the **no nsr** form of this command.

nsr
no nsr

| | |
|---------------------------|--|
| Syntax Description | This command has no arguments or keywords. |
| Command Default | BGP NSR is not activated. |
| Command Modes | XR Config |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **nsr** command to enable the Border Gateway Protocol (BGP) Nonstop Routing (NSR) with Stateful Switchover (SSO). This enables all **bgp** peerings to maintain the BGP state to ensure continuous packet forwarding during events that could interrupt service.



Note From release 5.2.3, NSR is enabled by default.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to enable BGP NSR:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 120
RP/0/RP0/CPU0:router(config-bgp)# nsr
```

The following example shows how to disable BGP NSR:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 120
RP/0/RP0/CPU0:router(config-bgp)# no nsr
```

Related Commands

| Command | Description |
|---|--|
| router bgp, on page 182 | Configures the Border Gateway Protocol (BGP) routing process. |
| nsr process-failures switchover | Configures failover as a recovery action in case of process failures for active instances to switch over to a standby route processor (RP) or a standby distributed route processor (DRP) to maintain nonstop routing (NSR). |
| show bgp nsr, on page 277 | Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information. |

nsr disable (BGP)

To disable Border Gateway Protocol (BGP) nonstop routing (NSR), use the **nsr disable** command in BGP XR Config mode. To re-enable BGP NSR, use the **no** form of this command.

nsr disable
no nsr disable

| | |
|---------------------------|--|
| Syntax Description | This command has no arguments or keywords. |
| Command Default | BGP NSR is activated by default. |
| Command Modes | XR Config |

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.3 | This command was introduced. |

Usage Guidelines Use the **nsr disable** command to disable Border Gateway Protocol (BGP) Nonstop Routing (NSR) with Stateful Switchover (SSO). Using the **no** form of this command enables all BGP peerings to maintain the BGP state to ensure continuous packet forwarding during events that could interrupt service.



Note In releases prior to R 5.2.3, NSR is disabled by default, and must be configured manually.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples The following example shows how to disable BGP NSR:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 120
RP/0/RP0/CPU0:router(config-bgp)# nsr disable
```

The following example shows how to re-enable BGP NSR:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# router bgp 120
RP/0/RP0/CPU0:router(config-bgp)# no nsr disable
```

Related Commands

| Command | Description |
|---|--|
| router bgp, on page 182 | Configures the Border Gateway Protocol (BGP) routing process. |
| nsr process-failures switchover | Configures failover as a recovery action in case of process failures for active instances to switch over to a standby route processor (RP) or a standby distributed route processor (DRP) to maintain nonstop routing (NSR). |
| show bgp nsr, on page 277 | Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information. |

orf

To specify Outbound Route Filter (ORF) and inbound filtering criteria, use the **orf route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

orf route-policy *route-policy-name*
no orf route-policy *route-policy-name*

| | | |
|---------------------------|--|------------------------------|
| Syntax Description | <i>route-policy-name</i> Name of the route policy. | |
| Command Default | No ORF route policy is defined. | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to configure outbound and inbound filtering criteria:

```
RP/0/RP0/CPU0:router(config)#router bgp 6
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)#orf route-policy policy_A
```

Related Commands

| Command | Description |
|---|--|
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |

password (BGP)

To enable Message Digest 5 (MD5) authentication on a TCP connection between two Border Gateway Protocol (BGP) neighbors, use the **password** command in an appropriate configuration mode. To disable MD5 authentication, use the **no** form of this command.

```
password {clear | encrypted} password
no password [{clear password | encrypted password}]
```

| Syntax Description | <table> <tr> <td data-bbox="386 569 503 600">clear</td><td data-bbox="516 569 1526 636">Specifies that an unencrypted password follows. The password must be a case-sensitive, clear-text unencrypted password.</td></tr> <tr> <td data-bbox="386 663 503 695">encrypted</td><td data-bbox="516 663 1526 730">Specifies that an encrypted password follows. The password must be a case-sensitive, encrypted password.</td></tr> <tr> <td data-bbox="386 751 503 783"><i>password</i></td><td data-bbox="516 751 1526 856">Password of up to 80 characters. The password can contain any alphanumeric characters. However, if the first character is a number or the password contains a space, the password must be enclosed in double quotation marks; for example, "2 password."</td></tr> </table> | clear | Specifies that an unencrypted password follows. The password must be a case-sensitive, clear-text unencrypted password. | encrypted | Specifies that an encrypted password follows. The password must be a case-sensitive, encrypted password. | <i>password</i> | Password of up to 80 characters. The password can contain any alphanumeric characters. However, if the first character is a number or the password contains a space, the password must be enclosed in double quotation marks; for example, "2 password." |
|---------------------------|---|--------------|---|------------------|--|-----------------|--|
| clear | Specifies that an unencrypted password follows. The password must be a case-sensitive, clear-text unencrypted password. | | | | | | |
| encrypted | Specifies that an encrypted password follows. The password must be a case-sensitive, encrypted password. | | | | | | |
| <i>password</i> | Password of up to 80 characters. The password can contain any alphanumeric characters. However, if the first character is a number or the password contains a space, the password must be enclosed in double quotation marks; for example, "2 password." | | | | | | |
| Command Default | When this command is not specified in the appropriate configuration mode, MD5 authentication is not enabled on a TCP connection between two BGP neighbors. | | | | | | |
| Command Modes | <p>Neighbor configuration</p> <p>Neighbor group configuration</p> <p>Session group configuration</p> | | | | | | |
| Command History | <table> <tr> <th data-bbox="386 1171 933 1203">Release</th><th data-bbox="959 1171 1526 1203">Modification</th></tr> <tr> <td data-bbox="386 1224 933 1255">Release 5.0.0</td><td data-bbox="959 1224 1526 1255">This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. | | |
| Release | Modification | | | | | | |
| Release 5.0.0 | This command was introduced. | | | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Configure a password to enable authentication between two BGP peers. Use the password command to verify each segment sent on the TCP connection between the peers. The same password must be configured on both networking devices, otherwise a connection cannot be made. The authentication feature uses the MD5 algorithm. Specifying this command causes the software to generate and check the MD5 digest on every segment sent on the TCP connection.</p> <p>Configuring a neighbor password does not cause the existing session for a neighbor to end. However, until the new password is configured on the remote router, the local BGP process does not receive keepalive messages from the remote device. If the password is not updated on the remote device by the end of the hold time, the session ends. The hold time can be changed using the timers command or the timers bgp command.</p> <p>If this command is configured for a neighbor group or neighbor address family group, a neighbor using the group inherits the configuration. Values of commands configured specifically for a neighbor overrides inherited values.</p> | | | | | | |

Task ID

Task ID Operations

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to configure neighbor 172.20.1.1 to use MD5 authentication with the password password1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)#password clear password1
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| password-disable, on page 163 | Overrides any inherited password configuration from a neighbor group or session group for BGP neighbors. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| timers (BGP), on page 351 | Set the timers for a specific BGP neighbor. |

password-disable

To override any inherited password configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **password-disable** command in an appropriate configuration mode. To disable overriding any inherited password command, use the **no** form of this command.

password-disable
no password-disable

| | | |
|--------------------|--|------------------------------|
| Syntax Description | This command has no arguments or keywords. | |
| Command Default | Configured passwords for neighbor and session groups are inherited. | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. If you specify a password on a neighbor group or session group, all users of the group inherit the password. Specifying a different password command specifically on a neighbor that uses the group overrides the inherited value. Specifying password-disable on a neighbor that uses the group disables password authentication for the neighbor. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to disable MD5 authentication for neighbor 172.20.1.1, preventing it from inheriting the password password1 from session group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# password clear password1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# password-disable
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| password (BGP), on page 161 | Enables MD5 authentication on a TCP connection between two BGP neighbors. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| use, on page 368 | Inherits characteristics from a neighbor group, a session group, or an address family group. |

precedence

To set the precedence level, use the **precedence** command in the appropriate configuration mode. To remove the **precedence** command from the configuration file and restore the system to its default interval values, use the **no** form of this command.

precedence *value*
no precedence [*value*]

| | | |
|--------------------|--|------------------------------|
| Syntax Description | <i>value</i> Value of the precedence. The precedence value can be a number from 0 to 7, or it can be one of the following keywords: critical —Set packets with critical precedence (5) flash — Set packets with flash precedence (3) flash-override —Set packets with flash override precedence (4) immediate —Set packets with immediate precedence (2) internet —Set packets with internetwork control precedence (6) network —Set packets with network control precedence (7) priority —Set packets with priority precedence (1) routine —Set packets with routine precedence (0) | |
| Command Default | No default behavior or values | |
| Command Modes | Neighbor configuration Neighbor session group configuration Neighbor group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. Use the precedence command to set the precedence value. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples


The following example shows how to set the precedence to 2:

```
RP/0/RP0/CPU0:router(config)# router bgp 5  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.1.1.1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100  
RP/0/RP0/CPU0:router(config-bgp-nbr)# precedence 2
```

receive-buffer-size

To set the size of the receive buffers for a Border Gateway Protocol (BGP) neighbor, use the **receive-buffer-size** command in an appropriate configuration mode. To remove the **receive-buffer-size** command from the configuration file and restore the system to its default condition in which the software uses the default size, use the **no** form of this command.

receive-buffer-size *socket-size* [*bgp-size*]
no receive-buffer-size [*socket-size*] [*bgp-size*]

| | | |
|--|--|------------------------------|
| Syntax Description | <i>socket-size</i> Size, in bytes, of the receive-side socket buffer. Range is 512 to 131072. | |
| | <i>bgp-size</i> (Optional) Size, in bytes, of the receive buffer in BGP. Range is 512 to 131072. | |
| Command Default | <i>socket-size</i> : 32,768 bytes <i>bgp-size</i> : 4,032 bytes | |
| Command Modes | Neighbor configuration | |
| | Neighbor group configuration | |
| | Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the receive-buffer-size command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because it allows the software to process a larger number of packets simultaneously. However, allocating larger buffers consumes more memory on the router. | |
|  Note | Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely. | |
| | If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to set the receive buffer sizes for neighbor 172.20.1.1 to be 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# receive-buffer-size 65536 8192
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-buffer-size, on page 183 | Sets the size of the send buffers for a BGP neighbor. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| socket receive-buffer-size, on page 341 | Sets the size of the receive buffers for all BGP neighbors. |

redistribute (BGP)

To redistribute routes from one routing domain into Border Gateway Protocol (BGP), use the **redistribute** command in an appropriate configuration mode. To disable route redistribution, use the **no** form of this command.

Connected

```
redistribute connected [metric metric-value] [route-policy route-policy-name]  
no redistribute connected [metric metric-value] [route-policy route-policy-name]
```

Enhanced Interior Gateway Routing Protocol (EIGRP)

```
redistribute eigrp process-id [match {external | internal}] [metric metric-value] [route-policy  
route-policy-name]  
no redistribute eigrp process-id [match {external | internal}] [metric metric-value] [route-policy  
route-policy-name]
```

Intermediate System-to-Intermediate System (IS-IS)

```
redistribute isis process-id [{level | {1 | 1-inter-area | 2}}] [metric metric-value] [route-policy  
route-policy-name]  
no redistribute isis process-id [{level | {1 | 1-inter-area | 2}}] [metric metric-value] [route-policy  
route-policy-name]
```

Open Shortest Path First (OSPF)

```
redistribute ospf process-id  
no redistribute ospf process-id
```

Routing Information Protocol

```
redistribute rip [metric metric-value] [route-policy route-policy-name]  
no redistribute rip [metric metric-value] [route-policy route-policy-name]
```

Static

```
redistribute static [metric metric-value] [route-policy route-policy-name]  
no redistribute static [metric metric-value] [route-policy route-policy-name]
```

| Syntax Description | | |
|--|--|--|
| connected | | Redistributes connected routes. Connected routes are established automatically when IP is enabled on an interface. |
| metric <i>metric-value</i> | | (Optional) Specifies the Multi Exit Discriminator (MED) attribute used for the redistributed route. Range is 0 to 4294967295. Use a value consistent with the destination protocol. By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0. |
| route-policy <i>route-policy-name</i> | | (Optional) Specifies a configured routing policy to filter redistributed routes. A route policy is used to filter the importation of routes from this source routing protocol to BGP. |
| eigrp | | Specifies that routes are distributed from EIGRP. You must be in IPv4 unicast address family configuration mode. |

| | |
|---|---|
| <i>process-id</i> | <p>For the eigrp keyword, an EIGRP instance name from which routes are to be redistributed.</p> <p>For the isis keyword, an IS-IS instance name from which routes are to be redistributed.</p> <p>For the ospf keyword, an OSPF instance name from which routes are to be redistributed.</p> <p>The <i>process-id</i> value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> |
| match { internal external [1 2] nssa-external [1 2] } | <p>(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:</p> <ul style="list-style-type: none"> • internal —Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes). • external [1 2]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • nssa-external [1 2]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. <p>For the external and nssa-external options, if a type is not specified, then both Type 1 and Type 2 are assumed.</p> |
| isis | <p>Specifies that routes are distributed from the IS-IS protocol.</p> <p>Redistribution from IS-IS is allowed under IPv4 unicast, IPv6 unicast, and address-families.</p> |
| level { 1 1-inter-area 2 } | <p>(Optional) Specifies the IS-IS level from which routes are redistributed. It can be one of the following:</p> <ul style="list-style-type: none"> • 1 —Routes are redistributed from Level 1 routes. • 1-inter-area —Routes are redistributed from Level 1 interarea routes. • 2 —Routes are redistributed from Level 2 routes. |
| ospf | Specifies that routes are distributed from the OSPF protocol. You must be in IPv4 unicast address family configuration mode. |
| rip | Specifies that routes are distributed from RIP. You must be in IPv4 unicast address family configuration mode. |
| static | Redistributes IP static routes. |

Command Default

Route redistribution is disabled.

For IS-IS, the default is to redistribute Level 1 and Level 2 routes.

For OSPF, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2.

For OSPFv3, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2.

By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0.

metric *metric-value*: 0

match { **internal** | **external** [1 | 2] | **nssa-external** [1 | 2]}: If no match is specified, the default is to match all routes.

Command Modes IPv4 address family configuration (**connected**, **eigrp**, **isis**, **ospf**, **rip**, and **static** are supported)

IPv6 address family configuration (**connected**, **eigrp**, **isis**,

ospfv3,
and **static** are supported)

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note When redistributing routes (into BGP) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Each instance of a protocol may be redistributed independently of the others. Changing or removing redistribution for a particular instance does not affect the redistribution capability of other protocols or other instances of the same protocol.

Networks specified using the **network** command are not affected by the **redistribute** command; that is, the routing policy specified in the **network** command takes precedence over the policy specified through the **redistribute** command.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples The following example shows how to redistribute IP Version 4 (IPv4) unicast OSPF routes from OSPF instance 110 into BGP:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# redistribute ospf 110
```

Related Commands

| Command | Description |
|--|---|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |

remote-as (BGP)

To create a Border Gateway Protocol (BGP) neighbor and begin the exchange of routing information, use the **remote-as** command in an appropriate configuration mode. To delete the entry for the BGP neighbor, use the **no** form of this command.

remote-as *as-number*
no remote-as [*as-number*]

Syntax Description

as-number Autonomous system (AS) to which the neighbor belongs.

- Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
- Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
- Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535.

Command Default

No BGP neighbors exist.

Command Modes

Neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **remote-as** command to create a neighbor and assign it a remote autonomous system number. A neighbor must have a remote autonomous system number before any other commands can be configured for it. Removing the remote autonomous system from a neighbor causes the neighbor to be deleted. You cannot remove the autonomous system number if the neighbor has other configuration.



Note

We recommend that you use the **no neighbor** command rather than the **no remote-as** command to delete a neighbor.

A neighbor specified with a remote autonomous system number that matches the autonomous system number specified in the **router bgp** command identifies the neighbor as internal to the local autonomous system. Otherwise, the neighbor is considered external.

Configuration of the **remote-as** command for a neighbor group or session group using the **neighbor-group** command or **session-group** command causes all neighbors using the group to inherit the characteristics configured with the command. Configuring the command directly for the neighbor overrides the value inherited from the group.

In the neighbor configuration submenu, configuring use of a session group or neighbor group for which **remote-as** is configured creates a neighbor and assigns it an autonomous system number if the neighbor has not already been created.



Note Do not combine **remote-as** commands and **no use neighbor-group** commands, or **remote-as** commands and **no use session-group** commands, in the same configuration commit.

Task ID

Task ID Operations

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to assign autonomous system numbers on two neighbors, neighbor 10.0.0.1, (internal) and neighbor 192.168.0.1 (external), setting up a peering session that shares routing information between this router and each of these neighbors:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group2
```

The following example shows how to configure a session group called group2 with an autonomous system number 1. Neighbor 10.0.0.1 is created when it inherits the autonomous system number 1 from session group group2.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group2
```

Related Commands

| Command | Description |
|---|--|
| neighbor (BGP), on page 139 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| router bgp, on page 182 | Configures the BGP routing process. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

| Command | Description |
|----------------------------------|---|
| use, on page 368 | Inherits characteristics from a neighbor group, session group, or address family group. |

remove-private-as

To remove private autonomous system numbers from autonomous system paths when generating updates to external neighbors, use the **remove-private-as** command in an appropriate configuration mode. To place the router in the default state in which it does not remove private autonomous system numbers, use the **no** form of this command.

remove-private-as [**inheritance-disable**] [**entire-aspath**]
no remove-private-as [**inheritance-disable**] [**entire-aspath**]

| Syntax Description | inheritance-disable (Optional) Permits the feature to be disabled from a neighbor group or address family group instead of being inherited. | | | | |
|---------------------------|--|---------|--------------|---------------|------------------------------|
| | entire-aspath (Optional) Removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private. | | | | |
| Command Default | When this command is not specified in the appropriate configuration mode, private autonomous system numbers are not removed from updates sent to external neighbors. | | | | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration | | | | |
| Command History | <table> <tr> <th data-bbox="347 1182 889 1234">Release</th><th data-bbox="889 1182 1489 1234">Modification</th></tr> <tr> <td data-bbox="347 1234 889 1295">Release 5.0.0</td><td data-bbox="889 1234 1489 1295">This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>This feature is available for external BGP (eBGP) neighbors only.</p> <p>When an update is passed to the external neighbor, the system drops any private autonomous system numbers. This happens irrespective of whether the autonomous system numbers are at the beginning or in the middle of the AS_SEQUENCE.</p> <p>If this command is used in a BGP confederation, the element following the confederation portion of the autonomous system path, if a sequence, is considered the leading sequence.</p> <p>The private autonomous system values range from 64512 to 65535.</p> <p>If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.</p> <p>Use the entire-aspath to removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private.</p> | | | | |

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows a configuration that removes the private autonomous system number from the IP Version 4 (IPv4) unicast updates sent to 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# remove-private-as
```

The following example shows how to disable the remove private autonomous system number feature for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# remove-private-as
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# remove-private-as inheritance-disable
```

Related Commands

| Command | Description |
|--|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| remote-as (BGP), on page 173 | Allows entries to the BGP neighbor table. |

route-policy (BGP)

To apply a routing policy to updates advertised to or received from a Border Gateway Protocol (BGP) neighbor, use the **route-policy** command in an appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *route-policy-name* [{*parameter1*, *parameter2*, . . . , *parametern*}] {**in** | **out**}
no route-policy *route-policy-name* [{*parameter1*, *parameter2*, . . . , *parametern*}] {**in** | **out**}

Syntax Description

route-policy-name Name of route policy. Up to 16 parameters can follow the route-policy-name, enclosed in brackets ([]).

in Applies policy to inbound routes.

out Applies policy to outbound routes.

Command Default

No policy is applied.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to specify a routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes. The **route-policy** command is used to define a policy.



Note

Configuring a large number of uniquely named outbound neighbor policies can adversely affect performance. This is true even if the uniquely named route policies are functionally identical. The user is discouraged from configuring multiple functionally identical route policies for use with this command. For example, if Policy A and Policy B are identical but named for different neighbors, the two policies should be configured as a single policy.

If the **route-policy** command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to apply the In-Ipv4 policy to inbound IP Version 4 (IPv4) unicast routes from neighbor 172.20.1.1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1  
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast  
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-policy In-Ipv4 in
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |

route-reflector-client

To configure the router as a Border Gateway Protocol (BGP) route reflector and configure the specified neighbor as its client, use the **route-reflector-client** command in an appropriate configuration mode. To disable configuring the neighbor as a client, use the **no** form of this command.

route-reflector-client [**inheritance-disable**]
no route-reflector-client [**inheritance-disable**]

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | inheritance-disable (Optional) Allows the configuration inherited from a neighbor group or address family group to be overridden. | |
| Command Default | The neighbor is not treated as a route reflector client. | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>This command is restricted to internal BGP (iBGP) neighbors only.</p> <p>Use the route-reflector-client command to configure the local router as the route reflector and the specified neighbor as one of its clients. All neighbors configured with this command are members of the client group, and the remaining iBGP peers are members of the nonclient group for the local route reflector.</p> <p>By default, all iBGP speakers in an autonomous system must be fully meshed with each other, and neighbors do not readvertise iBGP learned routes to other iBGP neighbors.</p> <p>With route reflection, all iBGP speakers need not be fully meshed. An iBGP speaker, the route reflector, passes learned iBGP routes to some number of iBGP client neighbors. Learned iBGP routes eliminate the need for each router running BGP to communicate with every other device running BGP in the autonomous system.</p> <p>The local router is a route reflector as long as it has at least one route reflector client.</p> <p>If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.</p> | |

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows neighbor at 172.20.1.1 configured as a route reflector client for IP Version 4 (IPv4) unicast routes:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

The following example disables the route-reflector client for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group1:

```
RP/0/RP0/CPU0:router(config)# router bgp 140
RP/0/RP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-afgrp)# route-reflector-client
RP/0/RP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client inheritance-disable
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| bgp cluster-id, on page 46 | Configures the cluster ID if the BGP cluster has more than one route reflector. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |

router bgp

To configure the Border Gateway Protocol (BGP) routing process, use the **router bgp** command in XR Config mode. To remove all BGP configurations and terminate the BGP routing process, use the **no** form of this command.

router bgp *as-number*

Syntax Description

as-number Number that identifies the autonomous system (AS) in which the router resides.

- Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
- Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
- Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535.

Command Default

No BGP routing process is enabled.

Command Modes

XR Config mode

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

Use the **router bgp** command to set up a distributed routing core that automatically guarantees the loop-free exchange of routing information between autonomous systems.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |
| rib | read, write |

Examples


The following example shows how to configure a BGP process for autonomous system 120:

```
RP/0/RP0/CPU0:router(config)# router bgp 120
```

send-buffer-size

To set the size of the send buffers for a Border Gateway Protocol (BGP) neighbor, use the **send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default values, use the **no** form of this command.

send-buffer-size *socket-size* [{*bgp-size*}]
no send-buffer-size [{*socket-size*}] [{*bgp-size*}]

| | | |
|--|---|------------------------------|
| Syntax Description | <i>socket-size</i> Size, in bytes, of the send-side socket buffer. Range is 4096 to 131072. | |
| | <i>bgp-size</i> (Optional) Size, in bytes, of the BGP process send buffer. Range is 4096 to 131072. | |
| Command Default | <i>socket-size</i> : 10240 bytes | |
| | <i>bgp-size</i> : 4096 bytes | |
| Command Modes | Use the socket send-buffer-size command to change the defaults. | |
| | Neighbor configuration | |
| | Neighbor group configuration | |
| | Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the send-buffer-size command to increase the buffer size employed when sending updates to a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on the router. | |
|  Note | Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses more memory indefinitely. | |
| | If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to set the send buffer sizes for neighbor 172.20.1.1 to be 8192 bytes for both the socket buffer and the BGP buffer:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# send-buffer-size 8192 8192
```

Related Commands

| Command | Description |
|--|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| receive-buffer-size, on page 167 | Sets the size of the receive buffers for a BGP neighbor. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| socket send-buffer-size, on page 343 | Sets the size of the send buffers for all BGP neighbors. |

send-community-ebgp

To specify that community or the large community attributes should be sent to an external Border Gateway Protocol (eBGP) neighbor, use the **send-community-ebgp** command in an appropriate configuration mode. To disable sending community or large community attributes to an eBGP neighbor, use the **no** form of this command.

```
send-community-ebgp [{inheritance-disable}]  
no send-community-ebgp [{inheritance-disable}]
```

| | | |
|---------------------------|--|--|
| Syntax Description | inheritance-disable (Optional) Allows configuration inherited from a neighbor group or address family group to be overridden. | |
| Command Default | Community (COMM) attributes are NOT sent to eBGP peers (including PE-CE peers). | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| | Release 6.3.1 | This command supports large communities. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the send-community-ebgp command to control whether community or large community attributes are sent to eBGP neighbors. This command cannot be configured for iBGP neighbors as community or large community attributes are always sent to iBGP neighbors.</p> <p>If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Configuring the command specifically for a neighbor overrides inherited values.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-extended-community-ebgp, on page 187 | Specifies that extended community attributes are sent to eBGP neighbors. |

send-extended-community-ebgp

To specify that extended community attributes should be sent to external Border Gateway Protocol (eBGP) neighbors, use the **send-extended-community-ebgp** command in an appropriate configuration mode. To disable sending extended community attributes to eBGP neighbors, use the **no** form of this command.

```
send-extended-community-ebgp [{inheritance-disable}]  
no send-extended-community-ebgp [{inheritance-disable}]
```

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | inheritance-disable (Optional) Allows configurations inherited from a neighbor group or address family group to be overridden. | |
| Command Default | Extended community (EXTCOMM) attributes are NOT sent to eBGP peers (including PE-CE peers). | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the send-extended-community-ebgp command to control whether extended community attributes are sent to eBGP neighbors. This command cannot be used for iBGP neighbors as extended community attributes are always sent to iBGP neighbors.</p> <p>If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Related Commands

| Command | Description |
|--|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-community-ebgp, on page 185 | Specifies that community attributes should be sent to an eBGP neighbor. |

session-group

To create a session group and enter session group configuration mode, use the **session-group** command in router configuration mode. To remove a session group and delete all configurations associated with it, use the **no** form of this command.

session-group *name*
no session-group *name*

| | |
|---------------------------|--|
| Syntax Description | name Name of the session group. |
|---------------------------|--|

| | |
|------------------------|--------------------------------|
| Command Default | No session groups are created. |
|------------------------|--------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

Use the **session-group** command to create a session group from which neighbors can inherit configuration that is address family-independent. That is, session groups cannot have address family-specific configuration. This command enters the session group configuration mode in which configuration for a session group is entered.

Many commands can be configured in both session group configuration mode and neighbor configuration mode.

Use of session groups saves time and reduces the router configuration size. Because the configuration of a session group can be inherited by any number of neighbors, use of the group can eliminate the need to copy long or complex configurations on each of a large number of neighbors. A neighbor can inherit all configuration from a session group simply by configuring the **use** command. Specific inherited session group configuration commands can be overridden for a specific neighbor by explicitly configuring the command for the specific neighbor.

The **no** form of this command causes all of the configuration for the session group to be removed. You cannot use the **no** form of this command if removing the group would leave one or more neighbors without a configured remote autonomous system number.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows a session group called group1 that is used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group.

```
RP/0/RSP0RP0/CPU0:router(config)# router bgp 1
RP/0/RSP0RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
```

The following example shows a session group called group1 used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group. However, the **password password1** configuration from group1 is overridden for neighbor 10.0.0.2, using the **password-disable** command in the neighbor 10.0.0.2 configuration submode.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# password password1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# password-disable
```

session-open-mode

To establish a Border Gateway Protocol (BGP) session with a specific TCP open mode, use the **session-open-mode** command in an appropriate configuration mode. To restore the default state, use the **no** form of this command.

session-open-mode {**active-only** | **both** | **passive-only**}
no session-open-mode [{**active-only** | **both** | **passive-only**}]

| | | |
|---------------------------|--|---|
| Syntax Description | active-only | Ensures that the BGP session can be established only when the request is initiated by the local end (active-open request) and all passive-open requests (from the other end) are rejected by the local BGP. |
| | both | Allows BGP sessions to be established from both incoming or outgoing TCP connection requests, with one being rejected in the event of a request collision. |
| | passive-only | Ensures that the local BGP does not initiate any TCP open requests and the session can be established only when the request comes from the remote end. |
| Command Default | The default is both . | |
| Command Modes | Neighbor configuration | |
| | Neighbor group configuration | |
| | Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | BGP, by default, tries to initiate an active TCP connection whenever a new neighbor is configured. A remote neighbor may also initiate the TCP connection before the local BGP can initiate the connection. This initiation of a TCP connection by a remote neighbor is considered a passive-open request and it is accepted by the local BGP. This default behavior can be modified using the session-open-mode command. | |

**Note**

The BGP connection is not opened and, as a result the BGP session, is not established if both the peering neighbors use the same nondefault TCP session open mode—active-only or passive-only. If both ends are configured with active-only, each neighbor rejects the TCP open request from the other end. One neighbor must be configured as passive-only or both. Similarly, if both neighbors are configured with passive-only, neither neighbor initiates the TCP open request and the BGP session is not established. Again, one neighbor must be configured as active-only or both. There is one exception. A connection open request from a neighbor that is configured with the TCP session open mode to be passive-only is processed to detect whether there is a connection collision before the request is rejected. This exception enables the local BGP to reset the session if the remote neighbor goes down and it is not detected by the local router.

Use the **session-open-mode** command when it may be necessary to preconfigure a neighbor that does not exist. Ensure that BGP does not spend any time actively trying to set up a TCP session with the neighbor. A BGP session does not come up between two neighbors, both of which configure the same nondefault value (**active-only** or **passive-only** keyword) for this command.

Task ID

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to enable a BGP session on router bgp 1:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 45.67.89.01
RP/0/RP0/CPU0:router(config-bgp-nbr)# session-open-mode active-only
```


show bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show bgp** command in XR EXEC mode.

show bgp [*ip-address* [{*mask* | /*prefix-length*}] [{**longer-prefixes** | **unknown-attributes** | **bestpath-compare**}]] [**standby**] [**detail**]

| Syntax | Description |
|---------------------------|--|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| <i>ip-address</i> | (Optional) Network address, entered to display a particular network in the BGP routing table. If the network address is omitted, then all networks in the BGP routing table are displayed. If the network mask and prefix length is omitted, then the software displays the longest matching prefix for the network address. |
| <i>mask</i> | (Optional) Network mask of the BGP route to match. |
| <i>/prefix-length</i> | (Optional) Prefix length of the BGP route to match. A slash (/) must precede the decimal value. |
| longer-prefixes | (Optional) Displays a route with the specified prefix length and more-specific routes if available. The longer-prefixes keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| unknown-attributes | (Optional) Includes unknown, transitive attributes. The unknown-attributes keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| bestpath-compare | (Optional) Displays route and best-path comparison information. The bestpath-compare keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| standby | (Optional) Displays information about the standby card. |
| detail | (Optional) Displays the prefix details. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. |
| Command Modes | XR EXEC mode |

Command History

| Release | Modification |
|---------------|--|
| Release 5.0.0 | This command was introduced. |
| Release 5.2.1 | The command output was modified to include information from rt-set configuration. |

Usage Guidelines

set default-afi



Note The **set default-afi** command is used to specify the default address family for the sessions and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for an address family or a subaddress family, each matching routing table is examined in turn.



Note Running the **show bgp** command immediately after configuring a large and complex route policy may result in timeout of the system database shown through an error message (SYSDB-SYSDB-6-TIMEOUT_EDM). It is recommended, that the show command be run, after the new route policy takes effect.

Use the **show bgp ip-address { mask | / prefix-length }** command to display detailed information for a specific route. If the mask and prefix length are omitted, the details of the longest matching prefix for the IP address are displayed.

Use the **show bgp** command to display all routes in the specified BGP routing table. Use the **show bgp ip-address { mask | / prefix-length } longer-prefixes** command to display those routes more specific than a particular prefix.

Use the **unknown-attributes** keyword to display details of any transitive attributes associated with a route that are not understood by the local system.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is the sample output from the **show bgp <IP address>** command displaying the graceful-shutdown community and the graceful-shut path attribute with BGP graceful maintenance feature activated:

```
RP/0/0/CPU0:R4#show bgp 5.5.5.5
...
10.10.10.1 from 10.10.10.1 (192.168.0.5)
```

```

Received Label 24000
Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best,
import-candidate
Received Path ID 0, Local Path ID 1, version 4
Community: graceful-shutdown
Originator: 192.168.0.5, Cluster list: 192.168.0.1
...

```

The following is sample output from the **show bgp** command in XR EXEC mode:

```

RP/0/RP0/CPU0:router#show bgp
BGP router identifier 172.20.1.1, local AS number 1820
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 3
Dampening enabled
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network          Next Hop          Metric LocPrf Weight Path
* i10.3.0.0/16    172.20.22.1        0      100      0 1800 1239 ?
*>i              172.20.16.1        0      100      0 1800 1239 ?
* i10.6.0.0/16    172.20.22.1        0      100      0 1800 690 568 ?
*>i              172.20.16.1        0      100      0 1800 690 568 ?
* i10.7.0.0/16    172.20.22.1        0      100      0 1800 701 35 ?
*>i              172.20.16.1        0      100      0 1800 701 35 ?
*                  192.168.40.24      0      100      0 1878 704 701 35 ?
* i10.8.0.0/16    172.20.22.1        0      100      0 1800 690 560 ?
*>i              172.20.16.1        0      100      0 1800 690 560 ?
*                  192.168.40.24      0      100      0 1878 704 701 560 ?
* i10.13.0.0/16   172.20.22.1        0      100      0 1800 690 200 ?
*>i              172.20.16.1        0      100      0 1800 690 200 ?
*                  192.168.40.24      0      100      0 1878 704 701 200 ?
* i10.15.0.0/16   172.20.22.1        0      100      0 1800 174 ?
*>i              172.20.16.1        0      100      0 1800 174 ?
* i10.16.0.0/16   172.20.22.1        0      100      0 1800 701 i
*>i              172.20.16.1        0      100      0 1800 701 i
*                  192.168.40.24      0      100      0 1878 704 701 i

Processed 8 prefixes, 8 paths

```

This table describes the significant fields shown in the display.

Table 2: show bgp Field Descriptions

| Field | Description |
|-----------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |

| Field | Description |
|--------------------------------|--|
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP table state | State of the BGP database. |
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between BGP scans for the specified address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |

| Field | Description |
|--------|---|
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the path origin code. |

The following is sample output from the **show bgp** command with the network specified:

```
RP/0/RP0/CPU0:router# show bgp 11.0.0.0/24
BGP router table entry for 11.0.0.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          2         2

Paths: (3 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    10.4.101.1
  Received by speaker 0
  Local
    0.0.0.0 from 0.0.0.0 (10.4.0.1)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, local, best
  Received by speaker 0
  2 3 4
    10.4.101.1 from 10.4.101.1 (10.4.101.1)
      Origin IGP, localpref 100, valid, external
  Received by speaker 0
  Local
    10.4.101.2 from 10.4.101.2 (10.4.101.2)
      Origin IGP, localpref 100, valid, internal
```

This table describes the significant fields shown in the display.

Table 3: show bgp prefix length Field Descriptions

| Field | Description |
|------------------------|--|
| BGP router table entry | Network that is being displayed. |
| Versions | List of the network versions in each BGP process. |
| Process | Name of the BGP process. |
| bRIB/RIB | Version of the network for sending to the RIB. You can compare this version with the bRIB/RIB version for the process (at the top of show bgp summary) to verify whether the network has been sent to the RIB. |

| Field | Description |
|---|---|
| SendTblVer | Version of the network for advertising to neighbors. This can be compared with the neighbor version to determine whether the network has been advertised to a particular neighbor. |
| Paths | List of paths for the network (that is, routes to reach the network). The number of paths and the index of the best path are given. |
| not advertised to any peer | Best path was received with a NO_ADVERTISE community and is not advertised to any neighbor. |
| not advertised to EBGp peer | Best path was received with a NO_EXPORT community and is not advertised to any eBGP neighbor. |
| not advertised outside local AS | Best path was received with a LOCAL_AS community and is not advertised to peers outside the local AS. |
| Advertisements of this net are suppressed by an aggregate | Network is a more-specific prefix of a configured aggregate and has been suppressed. It is not advertised to any neighbors unless they have an unsuppress-map configured. |
| Advertised to update-groups | List of update-groups to which the net has been advertised. Update-groups that have only one peer are not listed here. |
| Advertised to peers | List of neighbors to which the net has been advertised to. Neighbors that are in one of the update-groups listed above are not listed separately. Only neighbors that are in unique update-groups are listed. |
| Received by speaker 0 | BGP process where the path originated. This is always “speaker 0” for standalone mode. It will be the speaker-id when BGP is in distributed mode. |
| AS Path | Autonomous system (AS) path that was received for the path. If the AS path is empty, then “Local” is displayed. This is the case for paths that are locally generated on this router or on a neighboring router within the same AS. |
| aggregated by | If the path is an aggregate, the router-id of the router that performed the aggregation. |
| suppressed due to dampening | Path has been suppressed due to the configured path dampening. |
| history entry | Path is withdrawn, but a copy is kept to store the dampening information. |
| Received from a RR-client | Path was received from a route reflector client. |
| received-only | If soft reconfiguration inbound is configured, the path was received but dropped by inbound policy, or was accepted and modified. In either event, the received-only value is a copy of the original, unmodified path. |
| received & used | If soft reconfiguration inbound is configured, the path was received and accepted by inbound policy, but not modified. |
| stale | Neighbor from which the path was received is down, and the path is kept and marked as stale to support graceful restart. |

| Field | Description |
|--|--|
| <nexthop> from <neighbor> (<router-id>) | Next hop for the path. If the next hop is known by a mechanism outside BGP (for example, for redistributed paths), then 0.0.0.0 is displayed. After the next hop, the neighbor from whom the path was received is displayed, along with the neighbor's router-id. If the path was locally generated (for example, an aggregate or redistributed path), then 0.0.0.0 is displayed for the neighbor address. |
| Origin | IGP: the path originated from an IGP. EGP: the path originated from an EGP. incomplete: the origin of the path is unknown. |
| metric | MED value of the path. |
| localpref | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| weight | Locally assigned weight (if not 0) of the path. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| valid | Path is valid and can be considered in the best-path calculation. |
| redistributed | Path is redistributed through a redistribute command. |
| aggregated | Path is a locally generated aggregate created due to an aggregate-address command. |
| local | Path is a local network source due to a network command. |
| internal | Path was received from an iBGP neighbor. |
| external | Path was received from an eBGP neighbor. |
| atomic-aggregate | Path was received with the atomic-aggregate flag set. Some path information has been removed through aggregation. |
| best | Path is the best path for the network and is used for routing and advertised to peers. |
| multipath | Path is a multipath and is installed into the RIB along with the best path. |
| Community | List of communities attached to the path. |
| Extended community | List of extended communities attached to the path. |
| Originator | Originator of the path within the AS Cluster list if the path is reflected. |
| AS Cluster list | List of RR clusters the path has passed through if the path is reflected |
| Dampinfo | Penalty and reuse information if the path is dampened. |
| penalty | Current penalty for the path. |

| Field | Description |
|-----------------------|---|
| flapped | Number of times the path has flapped and the time since the first flap. |
| reuse in | Time until the path is re-used (undampened). |
| half life | Configured half-life for the path. |
| suppress value | Penalty at which the path is suppressed. |
| reuse value | Penalty at which the path is re-used. |
| Maximum suppress time | Maximum length of time for which the path can be suppressed. |

The following is sample output from the **show bgp** command with the *ip-address/prefix-length detail* options:

```
RP/0/RP0/CPU0:router# show bgp 51.0.0.0/24 detail

Sat Mar 14 00:37:14.109 PST PDT

BGP routing table entry for 51.0.0.0/24

Versions:

  Process          bRIB/RIB  SendTblVer

  Speaker          3         3

  Flags: 0x3e1000, label_retention: not enabled

Last Modified: Mar 13 19:32:17.976 for 05:04:56

Paths: (1 available, best #1)

  Advertised to update-groups (with more than one peer):

    0.3 0.4 0.7 0.8

  Advertised to peers (in unique update groups):

    201.48.20.1

  Path #1: Received by speaker 0

  Flags: 0x1000003

200 201

  213.0.0.6 from 213.0.0.6 (200.200.3.1)

  Origin IGP, localpref 100, valid, external, best
```

The following is sample output from the **show bgp** command with the additional paths received from:

```
BGP routing table entry for 51.0.1.0/24, Route Distinguisher: 2:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          63       63
```



```

Flags: 0x040630f2
Last Modified: Nov 11 12:44:05.811 for 00:00:16
Paths: (3 available, best #2)
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  Path #1: Received by speaker 0
  Flags: 0x3
  Not advertised to any peer
  111 111 111 111 111 111 111 111
    10.51.0.10 from 10.51.0.10 (11.11.11.11)
      Origin IGP, metric 0, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Extended community: RT:55:1
  Path #2: Received by speaker 0
  Flags: 0x5060007
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  561 562 563 564 565
    13.0.6.50 from 13.0.6.50 (13.0.6.50)
      Received Label 16
      Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported
      Received Path ID 0, Local Path ID 1, version 63
      Extended community: RT:55:1
  Path #3: Received by speaker 0
  Flags: 0x4060007
  Not advertised to any peer
  591 592 593 594 595
    13.0.9.50 from 13.0.9.50 (13.0.9.50)
      Received Label 16
      Origin IGP, localpref 100, valid, internal, backup, add-path, import-candidate,
imported
      Received Path ID 0, Local Path ID 4, version 63
      Extended community: RT:22:232 RT:55:1

```

The **show bgp** command output displays 'import suspect' when potential import oscillation has been detected for the prefix. Import of such a prefix is not affected. However, import of the prefix can be dampened in future if the oscillation continues. If the oscillation stops during the next import run, the prefix will no longer be marked 'import suspect'.

This is a sample output of **show bgp** command that displays large-communities in the list of attributes:

```

RP/0/RP0/CPU0:router#show bgp 3.3.3.3/32
Thu Mar 23 14:36:15.301 PDT
BGP routing table entry for 3.3.3.3/32
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          42         42
Last Modified: Mar 22 20:04:46.000 for 18:31:30
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    11.11.11.5
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    11.11.11.5
  Local
    10.10.10.3 from 10.10.10.3 (3.3.3.3)
      Origin incomplete, metric 0, localpref 94, valid, internal, best, group-best
      Received Path ID 0, Local Path ID 0, version 42
      Community: 258:259 260:261 262:263 264:265
      Large Community: 1:2:3 5:6:7 4123456789:4123456780:4123456788

```

Related Commands

| Command | Description |
|--|---|
| aggregate-address , on page 12 | Creates an aggregate entry in a BGP routing table. |
| bgp default local-preference , on page 54 | Changes the default local preference value. |
| network (BGP) , on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default subaddress Family Identifier (SAFI) for the current session. |
| show bgp cidr-only , on page 226 | Displays routes with nonnatural netmasks. |
| show bgp community , on page 230 | Displays routes belonging to the specified communities. |
| show bgp inconsistent-as , on page 247 | Displays networks with inconsistent origin autonomous system. |
| show bgp regexp , on page 308 | Displays routes matching an AS path regular expression. |
| show bgp route-policy , on page 312 | Displays networks that match a route policy. |
| show bgp summary , on page 320 | Displays the status of all BGP connections. |
| show bgp truncated-communities , on page 329 | Displays networks with community lists truncated by policy. |

show bgp large-community

To display routes that have the specified Border Gateway Protocol (BGP) large-communities, use the **show bgp large-community** command in XR EXEC mode.

```
show bgp [ ipv4 { unicast | labeled-unicast | all } ] large-community large-community-list
[ exact-match ]
show bgp [ ipv6 { unicast | labeled-unicast | all } ] large-community large-community-list
[ exact-match ]
show bgp [ all { unicast | labeled-unicast | all } ] large-community large-community-list [
exact-match ]
```

| Syntax Description | | |
|-----------------------------|---|---|
| ipv4 | (Optional) | Specifies IP Version 4 address prefixes. |
| unicast | (Optional) | Specifies unicast address prefixes. |
| labeled-unicast | (Optional) | Specifies labeled unicast address prefixes. |
| all | (Optional) | For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) | Specifies IP Version 6 address prefixes. |
| all | (Optional) | For address family, specifies prefixes for all address families. |
| large-community | Specifies that only routes with large-communities specified by <i>large-community-list</i> is displayed. | |
| <i>large-community-list</i> | They are specified as three non negative decimal integers separated by colons. For example, 1:2:3. Each integer is stored in 32 bits. The possible range for each integer is 0 to 4294967295. | |
| exact-match | (Optional) | Displays those routes that have large-communities exactly matching the specified large-communities. |

Command Default None.

Command Modes XR EXEC

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **exact-match** keyword to display only those routes with a set of large communities exactly matching the list of specified large communities. If you omit the **exact-match** keyword, those routes containing at least the specified large communities are displayed.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp large-community** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router#show bgp large-community 1:2:3 5:6:7
Thu Mar 23 14:40:33.597 PDT
BGP router identifier 4.4.4.4, local AS number 3
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000 RD version: 66
BGP main routing table version 66
BGP NSR Initial initsync version 3 (Reached)
BGP NSR/ISSU Sync-Group versions 66/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*>i3.3.3.3/32      10.10.10.3             0      94      0 ?
*> 5.5.5.5/32      11.11.11.5             0              0 5 ?
```

In the above example, the routes 3.3.3.3/32 and 5.5.5.5/32 contain both the large communities 1:2:3 and 5:6:7 and some additional large communities.

```
RP/0/RP0/CPU0:router#show bgp large-community 1:2:3 5:6:7 exact-match
Thu Mar 23 14:50:28.155 PDT
```

In the above example, no routes are displayed because the available routes have some additional large-communities, so the match is not exact.

Related Commands

| Command | Description |
|---|--|
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| show bgp, on page 193 | Displays BGP routes. |

send-community-ebgp

To specify that community or the large community attributes should be sent to an external Border Gateway Protocol (eBGP) neighbor, use the **send-community-ebgp** command in an appropriate configuration mode. To disable sending community or large community attributes to an eBGP neighbor, use the **no** form of this command.

```
send-community-ebgp [{inheritance-disable}]
no send-community-ebgp [{inheritance-disable}]
```

| | | |
|--------------------|--|--|
| Syntax Description | inheritance-disable (Optional) Allows configuration inherited from a neighbor group or address family group to be overridden. | |
| Command Default | Community (COMM) attributes are NOT sent to eBGP peers (including PE-CE peers). | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| | Release 6.3.1 | This command supports large communities. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the send-community-ebgp command to control whether community or large community attributes are sent to eBGP neighbors. This command cannot be configured for iBGP neighbors as community or large community attributes are always sent to iBGP neighbors.</p> <p>If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Configuring the command specifically for a neighbor overrides inherited values.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-extended-community-ebgp, on page 187 | Specifies that extended community attributes are sent to eBGP neighbors. |

show bgp update out

To display address-family level update generation information, use the **show bgp update out** command in XR EXEC mode.

```
show bgp [afi safi] update out [{brief | detail}]
```

| | | |
|--------------------|---|------------------------------|
| Syntax Description | <i>afi</i> (Optional) Displays address-family identifier. | |
| | <i>safi</i> (Optional) Displays subsequent address family identifier. | |
| | brief (Optional) Displays brief information on process level update generation. | |
| | detail (Optional) Displays detailed information on process level update generation. | |
| Command Default | None | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operation |
| | bgp | read |

show bgp update in error process

To display process level update inbound error-handling information, use the **show bgp update in error process** command in XR EXEC mode.

show bgp update in error process [{**brief** | **detail**}]

| | | |
|--------------------|---|--|
| Syntax Description | brief | (Optional) Displays brief information on process level update generation. |
| | detail | (Optional) Displays detailed information on process level update generation. |
| Command Default | None | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read |

This example displays sample output from the **show bgp update in error process** command:

```
RP/0/RP0/CPU0:router#show bgp update in error process
```

```
Basic Update error-handling:
```

```
EBGP: [Enabled]
```

```
IBGP: [Enabled]
```

```
Extended Update error-handling:
```

```
EBGP: [Disabled]
```

```
IBGP: [Disabled]
```

```
Malformed Update messages: 0
```

```
Neighbors that received malformed Update messages: 0
```

```
Last malformed Update received: --- (never)
```


show bgp update out filter-group

To display update generation information at filter-group level, **show bgp update out filter-group** command in XR EXEC mode.

```
show bgp [afi safi] update out filter-group [fg-process-id] [{brief | detail}]
```

| | | |
|--------------------|----------------------|--|
| Syntax Description | <i>afi safi</i> | Specifies the address family and subsequent address family identifiers. |
| | <i>fg-process-id</i> | Specifies the filter-group process ID in <x.y> format. Range is <0-15>.<0-4294967295>. |
| | brief | (Optional) Displays brief information on filter-group level update generation |
| | detail | (Optional) Displays detailed information on filter-group level update generation. |

| | |
|-----------------|------|
| Command Default | None |
|-----------------|------|

| | |
|---------------|---------|
| Command Modes | XR EXEC |
|---------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|------------------|---|

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

This example displays sample output from **show bgp update out filter-group** command:

show bgp update out process

To display process level update generation information, use the **show bgp update out process** command in XR EXEC mode.

show bgp update out process [{brief | detail}]

| | |
|---------------------------|---|
| Syntax Description | brief (Optional) Displays brief information on process level update generation. |
| | detail (Optional) Displays detailed information on process level update generation. |
| Command Default | None |
| Command Modes | XR EXEC |
| Command History | Release Modification |
| | Release This command was introduced. 5.0.0 |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
| Task ID | Task ID Operation |
| | bgp read |

This example displays sample output from the show bgp update out process brief command:

```
RP/0/RP0/CPU0:router#show bgp update out process
Wed Sep 12 08:26:04.308 DST

Update generation status: Normal
Update OutQ:                0 bytes (0 messages)
Update limit: 536870912 bytes (configured 536870912 bytes)

Update generation logging: [Disabled]

  Address-family Status   Limit      OutQ      UG   SG (Thr)   SG-R (Thr) Nbrs
-----
IPv4 Unicast   Normal   268435456   0           1    1(0)      0(0)       3

History:
Update OutQ Hi:                300 bytes (1 messages)
Update OutQ Cumulative:        1200 bytes (4 messages)
Update OutQ Discarded:          0 bytes (0 messages)
Update OutQ Cleared:            0 bytes (0 messages)
Last discarded from OutQ: --- (never)
Last cleared from OutQ: --- (never)
Update generation throttled 0 times, last event --- (never)
```

```
Update generation recovered 0 times, last event --- (never)
Update generation mem alloc failed 0 times, last event --- (never)
```

show bgp update out sub-group

To display sub-group update generation information, use the **show bgp update out sub-group** command in XR EXEC mode.

show bgp [*afi safi*] **update out** [**update-group** *ug-index*] **sub-group** [*sg-index*] [{**brief** | **detail**}]

Syntax Description

| | |
|-----------------|--|
| <i>afi</i> | (Optional) Displays address-family identifier. |
| <i>safi</i> | (Optional) Displays subsequent address family identifier. |
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |
| <i>ug-index</i> | (Optional) Displays the update-group process ID in <x.y> format. |
| <i>sg-index</i> | (Optional) displays the sub-group process ID in <x.y> format. |

Command Default

None

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |

show bgp update out update-group

To display update-group update generation information, use the **show bgp update out update-group** command in XR EXEC mode.

```
show bgp [afi safi] update out update-group [ug-index] [{brief | detail}]
```

| | | |
|--------------------|---|------------------------------|
| Syntax Description | <i>afi</i> (Optional) Displays address-family identifier. | |
| | <i>safi</i> (Optional) Displays subsequent address family identifier. | |
| | brief (Optional) Displays brief information on process level update generation. | |
| | detail (Optional) Displays detailed information on process level update generation. | |
| | <i>ug-index</i> (Optional) Displays the update-group process ID in <x.y> format. | |
| Command Default | None | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operation |
| | bgp | read |

show bgp advertised

To display advertisements for neighbors or a single neighbor, use the **show bgp advertised** command in XR EXEC mode.

```
show bgp [ipv4 { all | labeled-unicast | unicast }] advertised [neighbor ip-address] [standby]
[summary]
```

```
show bgp [ipv6 { all | labeled-unicast | unicast }] advertised [neighbor ip-address] [standby]
[summary]
```

```
show bgp [all { all | labeled-unicast | unicast }] advertised [neighbor ip-address] [standby]
[summary]
```

| | | |
|--------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | neighbor | (Optional) Previews advertisements for a single neighbor. If the neighbor keyword is omitted, then the advertisements for all neighbors are displayed. |
| | <i>ip-address</i> | (Optional) IP address of the neighbor. |
| | standby | (Optional) Displays information about the standby card. |
| | summary | (Optional) Displays a summary of advertisements. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that is configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp advertised** command to display the routes that have been advertised to peers or a specific peer. To preview advertisements that would be sent to a peer under a particular policy, even if the corresponding update messages have not been generated yet, use the **show bgp policy** command.



Note When you issue the **show bgp advertised** command, a route is not displayed in the output unless an advertisement for that route has already been sent (and not withdrawn). If an advertisement for the route has not yet been sent, the route is not displayed.

Use the **summary** keyword to display a summary of the advertised routes. If you do not specify the **summary** keyword, the software displays detailed information about the advertised routes.



Note The **show bgp advertised** command does not display the application of any outbound policy in the route details it displays. Consequently, this command provides only an indication of whether a particular route has been advertised, rather than details of which attributes were advertised. Use the **show bgp policy sent-advertisements** command to display the attributes that are advertised.

Task ID

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|-----|------|
| bgp | read |
|-----|------|

Examples

The following is sample output from the **show bgp advertised** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp advertised neighbor 10.0.101.4 summary
```

| Network | Next Hop | From | AS Path |
|------------|------------|------------|---------------------------|
| 1.1.1.0/24 | 10.0.101.1 | 10.0.101.1 | 2 3 222 333 444 555 i |
| 1.1.2.0/24 | 10.0.101.1 | 10.0.101.1 | 3 4 5 6 7 i |
| 1.1.3.0/24 | 10.0.101.1 | 10.0.101.1 | 77 88 33 44 55 99 99 99 i |
| 1.1.4.0/24 | 10.0.101.1 | 10.0.101.1 | 2 5 6 7 8 i |
| 1.1.7.0/24 | 10.0.101.1 | 10.0.101.1 | 3 5 i |
| 1.1.8.0/24 | 10.0.101.1 | 10.0.101.1 | 77 88 99 99 99 i |

This table describes the significant fields shown in the display.

Table 4: show bgp advertised neighbor summary Field Descriptions

| Field | Description |
|-----------------|--|
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| From | IP address of the peer that advertised this route. |
| AS Path | AS path of the peer that advertised this route. |
| Local | Indicates the route originated on the local system. |
| Local Aggregate | Indicates the route is an aggregate created on the local system. |
| Advertised to | Indicates the peer to which this entry was advertised. This field is used in the output when displaying a summary of the advertisements to all neighbors. |

The following is sample output from the **show bgp advertised** command for detailed advertisement information:

```
RP/0/RP0/CPU0:router# show bgp advertised neighbor 172.72.77.1

172.16.0.0/24 is advertised to 172.72.77.1
  Path info:
    neighbor: Local          neighbor router id: 172.74.84.1
    valid redistributed best
  Attributes after inbound policy was applied:
  next hop: 0.0.0.0
    MET ORG AS
    origin: incomplete metric: 0
    aspath:
10.52.0.0/16 is advertised to 172.72.77.1
  Path info:
    neighbor: Local Aggregate neighbor router id: 172.74.84.1
    valid aggregated best
  Attributes after inbound policy was applied:
  next hop: 0.0.0.0
    ORG AGG ATOM
    origin: IGP aggregator: 172.74.84.1 (1)
    aspath:
```

This table describes the significant fields shown in the display.

Table 5: show bgp advertised neighbor Field Descriptions

| Field | Description |
|------------------|--|
| is advertised to | IP address of the peer to which this route has been advertised. If the route has been advertised to multiple peers, the information is shown separately for each peer. |

| Field | Description |
|-------------------------------------|--|
| neighbor | IP address of the peer that advertised this route, or one of the following: Local—Route originated on the local system. Local Aggregate—Route is an aggregate created on the local system. |
| neighbor router id | BGP identifier for the peer, or the local system if the route originated on the local system. |
| Not advertised to any peer | Indicates the no-advertise well-known community is associated with this route. Routes with this community are not advertised to any BGP peers. |
| Not advertised to any EBGp peer | Indicates the no-export well-known community is associated with this route. Routes with this community are not advertised to external BGP peers, even if those external peers are part of the same confederation as the local router. |
| Not advertised outside the local AS | Indicates the local-AS well-known community is associated with this route. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. |
| (Received from a RR-client) | Path was received from a route reflector client. |
| (received-only) | This path is not used for routing purposes. It is used to support soft reconfiguration, and records the path attributes before inbound policy was applied to a path received from a peer. A path marked “received-only” indicates that either the path was dropped by inbound policy, or the path information was modified by inbound policy and a separate copy of the modified path is used for routing. |
| (received & used) | Indicates that the path is used both for soft reconfiguration and routing purposes. A path marked “received and used,” implies the path information was not modified by inbound policy. |
| valid | Path is valid. |
| redistributed | Path is locally sourced through redistribution. |
| aggregated | Path is locally sourced through aggregation. |
| local | Path is locally sourced through the network command. |
| confed | Path was received from a confederation peer. |
| best | Path is selected as best. |
| multipath | Path is one of multiple paths selected for load-sharing purposes. |

| Field | Description |
|---|--|
| dampinfo | Indicates dampening information: Penalty—Current penalty for this path. Flapped—Number of times the route has flapped. In—Time (hours:minutes:seconds) since the router noticed the first flap. Reuse in—Time (hours:minutes:seconds) after which the path is made available. This field is displayed only if the path is currently suppressed. |
| Attributes after inbound policy was applied | Displays attributes associated with the received route, after any inbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present. |
| origin | Origin of the path: IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command. EGP—Path originated from an Exterior Gateway Protocol. incomplete—Origin of the path is not clear. For example, a route that is redistributed into BGP from an IGP. |
| neighbor as | First autonomous system (AS) number in the AS path. |
| aggregator | Indicates that the path was received with the aggregator attribute. The autonomous system number and router-id of the system that performed the aggregation are shown. |
| metric | Value of the interautonomous system metric, otherwise known as the MED metric. |
| localpref | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| aspath | AS path associated with the route. |

| Field | Description |
|--------------------|---|
| community | <p>Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities:</p> <p>Local-AS—Community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise—Community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export—Community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation with the local router.</p> |
| Extended community | <p>Extended community attributes associated with the path. For known extended community types, the following codes may be displayed:</p> <p>RT—Route target community</p> <p>SoO—Site of Origin community</p> <p>LB—Link Bandwidth community</p> |
| Originator | Router ID of the originating router when route reflection is used. |
| Cluster lists | Router ID or cluster ID of all route reflectors through which the route has passed. |

Related Commands

| Command | Description |
|---|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default subaddress Family Identifier (SAFI) for the current session. |
| route-policy (BGP), on page 178 | Applies a route policy to incoming and outgoing routes. |
| show bgp policy, on page 283 | Displays information about BGP advertisements under a proposed policy. |
| sent-advertisements | Previews advertisements to peers, including details of advertised attributes. |

show bgp af-group

To display information about Border Gateway Protocol (BGP) configuration for address family groups, use the **show bgp af-group** command in XR EXEC mode.

show bgp af-group *group-name* {**configuration** [**defaults**] [**nvgen**] | **inheritance** | **users**}

Syntax Description

| | |
|----------------------|--|
| <i>group-name</i> | Name of the address family group to display. |
| configuration | (Optional) Displays the effective configuration for the af-group, including any settings that have been inherited from af-groups used by this af-group. |
| defaults | (Optional) Displays all configuration settings, including any default settings. |
| nvgen | (Optional) Displays output in the format of show running-config output. If the defaults keyword is also specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | Displays the af-groups from which this af-group inherits configuration settings. |
| users | Displays the neighbors, neighbor groups, and af-groups that inherit configuration from this af-group. |

Command Default

No default behavior or value

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp af-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of an af-group, taking into account any configuration that may be inherited from other af-groups through the **use af-group** command. The source of each command is shown.

If the **defaults** keyword is specified, all configuration for the af-group, including default values, is shown. Default configuration is identified in the show output. Use the **nvgen** keyword to display configuration formatted in the style of the **show running-config** command. This output is suitable for cutting and pasting into configuration sessions.

Use the **show bgp af-group** command with the *group-name* **inheritance** argument and keyword to display the address family groups from which the specified af-group inherits configuration.

Use the **show bgp af-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and af-groups that inherit configuration from the specified af-group.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following af-group configuration is used in the examples:

```
af-group group3 address-family ipv4 unicast
remove-private-AS
soft-reconfiguration inbound
!
af-group group1 address-family ipv4 unicast
use af-group group2
maximum-prefix 2500 75 warning-only
default-originate
soft-reconfiguration inbound disable
!
af-group group2 address-family ipv4 unicast
use af-group group3
send-community-ebgp
send-extended-community-ebgp
capability orf prefix both
```

The following is sample output from the **show bgp af-group** command with the **configuration** keyword in XR EXEC mode. The source of each command is shown in the right column. For example, **default-originate** is configured directly on **af-group group1**, and the **remove-private-AS** command is inherited from af-group group2, which in turn inherits it from af-group group3.

```
RP/0/RP0/CPU0:router# show bgp af-group group1 configuration

af-group group1 address-family ipv4 unicast
  capability orf prefix both           [a:group2]
  default-originate                    []
  maximum-prefix 2500 75 warning-only  []
  remove-private-AS                    [a:group2 a:group3]
  send-community                       [a:group2]
  send-extended-community               [a:group2]
```

The following is sample output from the **show bgp af-group** command with the **users** keyword:

```
RP/0/RP0/CPU0:router# show bgp af-group group2 users

IPv4 Unicast: a:group1
```

The following is sample output from the **show bgp af-group** command with the **inheritance** keyword. This example shows that the specified af-group group1 directly uses the group2 af-group, which in turn uses the group3 af-group:

```
RP/0/RSP0RP0/CPU0:router# show bgp af-group group1 inheritance

IPv4 Unicast: a:group2 a:group3
```

Table 6: [show bgp af-group Field Descriptions, on page 222](#) describes the significant fields shown in the display.

This table describes the significant fields shown in the display.

Table 6: show bgp af-group Field Descriptions

| Field | Description |
|-----------|--|
| [] | Configures the command directly on the specified address family group. |
| a: | Indicates the name that follows is an address family group. |
| n: | Indicates the name that follows is a neighbor group. |
| [dflt] | Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the defaults keyword is specified. |
| <not set> | Indicates that the configuration is disabled by default. This field may be shown when the defaults keyword is specified. |

Related Commands

| Command | Description |
|--|--|
| af-group, on page 10 | Configures a BGP address family group. |
| show bgp neighbors, on page 256 | Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups. |
| show bgp neighbor-group, on page 252 | Displays information about configuration for neighbor groups. |
| use, on page 368 af-group | Configures an af-group to inherit the configuration of a specified af-group. |

show bgp attribute-key

To display all existing attribute keys, use the **show bgp attribute-key** command in XR EXEC mode.

show bgp {ipv4 | ipv6 | all} attribute-key [standby]

| | | |
|--------------------|------------------------|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | all | (Optional) For subaddress family, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | standby | (Optional) Displays information about the standby card. |

| | |
|-----------------|--|
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. |
|-----------------|--|

| | |
|---------------|---------|
| Command Modes | XR EXEC |
|---------------|---------|

| | | |
|-----------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

| | |
|------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|------------------|---|



| | |
|------|--|
| Note | The set default-afi command is used to specify the default address family for the session, and the set default-safi command is used to specify the default subaddress family for the session. See the <i>System Management Command Reference for Cisco NCS 6000 Series Routers</i> for detailed information and syntax for the set default-afi and set default-safi commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast. |
|------|--|

| | | |
|---------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read |

| | |
|----------|--|
| Examples | The following is sample output from the show bgp attribute-key command in XR EXEC mode: |
|----------|--|

```

RP/0/RP0/CPU0:router# show bgp all all attribute-key

Address Family: IPv4 Unicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 109
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop        AttrKey
*> 1.1.0.0/16      0.0.0.0          0x00000002
*> 10.0.0.0/16     0.0.0.0          0x00000002
*> 12.21.0.0/16    0.0.0.0          0x00000002
*> 194.3.192.1/32  10.0.101.1       0x00000009
*> 194.3.192.2/32  10.0.101.1       0x00000009
*> 194.3.192.3/32  10.0.101.1       0x00000009
*> 194.3.192.4/32  10.0.101.1       0x00000009
*> 194.3.192.5/32  10.0.101.1       0x00000009

Processed 8 prefixes, 8 paths

Address Family: IPv6 Unicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 19
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop        AttrKey
*> 2222::1111/128  2222::2         0x00000009
*> 2222::1112/128  2222::2         0x00000009

Processed 2 prefixes, 2 paths

```

This table describes the significant fields shown in the display.

Table 7: show bgp attribute-key Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| BGP scan interval | Interval (in seconds) between scans. |

| Field | Description |
|---|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Entry originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| AttrKey | Key associated with the route attribute. |
| Processed <i>n</i> prefixes, <i>n</i> paths | Number of prefixes and number of paths processed for the table. |

Related Commands

| Command | Description |
|-------------------------|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |

show bgp cidr-only

To display routes with nonnatural network masks, also known as classless interdomain routing (CIDR) routes, use the **show bgp cidr-only** command in XR EXEC mode.

show bgp [{**ipv4**}] **cidr-only** [**standby**]

| | | |
|---------------------------|---|--|
| Syntax Description | ipv4 | (Optional) Specifies the IP Version 4 address family. |
| | unicast | (Optional) Specifies the unicast address family. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress family, specifies all subaddress families. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. This command is applicable only for IPv4 prefixes. If the default address family is not IPv4, then the ipv4 keyword must be used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for subaddress family, all subaddress family routing tables are examined.

The **show bgp cidr-only** command applies only for IPv4 prefixes. If the **ipv4** keyword is not specified and the default address family is not IPv4, the command is not available.

Use the **show bgp cidr-only** command to display CIDR routes. Routes that have their correct class (class A, B, or C) prefix length are not displayed.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp cidr-only** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp cidr-only

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 2589
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop        Metric    LocPrf    Weight    Path
*> 192.0.0.0/8    192.168.72.24      0         1878     ?
*> 192.168.0.0/16 192.168.72.30      0         108      ?
```

This table describes the significant fields shown in the display.

Table 8: show bgp cidr-only Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Entry originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|---|
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp, on page 193 | Displays BGP routes. |

show bgp community

To display routes that have the specified Border Gateway Protocol (BGP) communities, use the **show bgp community** command in XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] community community-list [exact-match]
show bgp [ipv6 {unicast | labeled-unicast | all}] community community-list [exact-match]
show bgp [all {unicast | labeled-unicast | all}] community community-list [exact-match]
```

| Syntax Description | |
|------------------------|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| community | Specifies that only routes with communities specified by <i>community-list</i> is displayed. |
| <i>community-list</i> | <p>Between one and seven communities. Each community can be a number in the range from 1 to 4294967295, a community specified in AA:NN format, or one of the following well-known communities:</p> <p>local-AS —Well-known community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise —Well-known community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export —Well-known community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.</p> <p>internet —Well-known community whose value is not defined in BGP RFC. IOS XR BGP uses a value of 0 for the internet community. Routes with this community are advertised to all peers without any restrictions.</p> <p>For the AA:NN format:</p> <p>AA—Range is 0 to 65535.</p> <p>NN—Range is 1 to 4294967295.</p> <p>Up to seven community numbers can be specified.</p> |
| exact-match | (Optional) Displays those routes that have communities exactly matching the specified communities. |

| | | |
|-------------------------|---|------------------------------|
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or the subaddress family, each matching routing table is examined in turn.

If more than seven communities are required, it is necessary to configure a route policy and use the [show bgp route-policy, on page 312](#) command.

Use the **exact-match** keyword to display only those routes with a set of communities exactly matching the list of specified communities. If you omit the **exact-match** keyword, those routes containing at least the specified communities are displayed.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

The following is sample output from the **show bgp community graceful-shutdown** command displaying the graceful maintenance feature information:

```
RP/0/0/CPU0:R4#show bgp community graceful-shutdown
Tue Jan 27 13:36:25.006 PST
BGP router identifier 192.168.0.4, local AS number 4
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 18
BGP main routing table version 18
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
               Network      Next Hop          Metric LocPrf Weight Path
```

show bgp community

```
* 5.5.5.5/32      10.10.10.1      88      0 1 ?

Processed 1 prefixes, 1 paths
```

Examples

The following is sample output from the **show bgp community** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp community 1820:1 exact-match

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 55
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop          Metric LocPrf Weight Path
*  10.13.0.0/16    192.168.40.24                    0 1878 704 701 200 ?
*  10.16.0.0/16    192.168.40.24                    0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 9: show bgp community Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|---|
| aggregate-address, on page 12 | Creates an aggregate entry in a BGP routing table. |
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp, on page 193 | Displays BGP routes. |

show bgp convergence

To display whether a specific address family has reached convergence, use the **show bgp convergence** command in XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] convergence
show bgp [ipv6 {unicast | labeled-unicast | all}] convergence
show bgp [all {unicast | labeled-unicast | all}] convergence
```

| | | |
|--------------------|---|--|
| Syntax Description | ipv4 | (Optional) Specifies the IP Version 4 address family. |
| | unicast | (Optional) Specifies the unicast address family. |
| | labeled-unicast | (Optional) Specifies unicast address prefixes. |
| | all | (Optional) For subaddress family, specifies all subaddress families. |
| | ipv6 | (Optional) Specifies the IP Version 6 address family. |
| | all | (Optional) For address family, specifies all address families. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp convergence** command to see if there is any pending work for BGP to perform. The software checks the following conditions to determine whether the specified address family has converged. If all the conditions are true, the address family is considered converged.

- All received updates have been processed and best routes selected.
- All selected routes have been installed in the global Routing Information Base (RIB).
- All selected routes have been advertised to peers, including any peers that are not established (unless those peers have been administratively shut down). See the **shutdown (BGP)** command for more information about administrative shutdown.

While testing that all selected routes have been advertised to peers, the **show bgp convergence** command checks the size of the write queue for each neighbor. Because this queue is shared by all address families, there is a small possibility that the command indicates the address family has not converged when, in fact, it has converged. This could happen if the neighbor write queue contained messages from some other address family.

If the specified address family has not converged, the **show bgp convergence** command output does not indicate the amount of work that is pending. To display this information, use the **show bgp summary** command.

Task ID

Task ID Operations

bgp read

Examples

The following shows the result of using the **show bgp convergence** command for an address family that has converged:

```
RP/0/RP0/CPU0:router# show bgp convergence
```

```
Converged.
All received routes in RIB, all neighbors updated.
All neighbors have empty write queues.
```

The following shows the result of using the **show bgp convergence** command for an address family that has not converged:

```
RP/0/RP0/CPU0:router# show bgp convergence
```

```
Not converged.
Received routes may not be entered in RIB.
One or more neighbors may need updating.
```

This table describes the significant fields shown in the display.

Table 10: show bgp convergence Field Descriptions

| Field | Description |
|-------------------------|---|
| Converged/Not converged | Specifies whether or not all routes have been installed in the RIB and updates have been generated and sent to all neighbors. |

| Field | Description |
|----------------------------------|--|
| [All] Received routes... | For convergence, all routes must have been installed into the RIB and all updates must have been generated. For non-convergence, some routes may not be installed in the RIB, or some routes that have been withdrawn have not yet been removed from the RIB, or some routes that are up to date in the RIB have not been advertised to all neighbors. |
| [All One or more] neighbors... | Specifies the status of neighbor updating. |

Related Commands

| Command | Description |
|---|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp summary, on page 320 | Displays the status of BGP peer connections. |
| shutdown (BGP), on page 339 | Disables a neighbor without removing all of its configuration. |

show bgp dampened-paths

To display Border Gateway Protocol (BGP) dampened routes, use the **show bgp dampened-paths** command in XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] dampened-paths [standby]
show bgp [ipv6 {unicast | labeled-unicast | all}] dampened-paths [standby]
```

| Syntax Description | | |
|------------------------|------------|--|
| ipv4 | (Optional) | Specifies IP Version 4 address prefixes. |
| unicast | (Optional) | Specifies unicast address prefixes. |
| labeled-unicast | (Optional) | Specifies labeled unicast address prefixes. |
| all | (Optional) | For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) | Specifies IP Version 6 address prefixes. |
| all | (Optional) | For address family, specifies prefixes for all address families. |
| standby | (Optional) | Displays information about the standby card. |

| | |
|------------------------|--|
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. |
|------------------------|--|

| | |
|----------------------|---------|
| Command Modes | XR EXEC |
|----------------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp dampened-paths** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp dampened-paths

BGP router identifier 10.2.0.1, local AS number 3
BGP generic scan interval 60 secs
BGP main routing table version 7
Dampening enabled
BGP scan interval 60 secs
Status codes:s suppressed, d damped, h history, * valid, > best
               i - internal, S stale

Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          From             Reuse      Path
*d 10.0.0.0         10.0.101.35    00:01:20  35 i
```

This table describes the significant fields shown in the display.

Table 11: show bgp dampened-paths Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| From | Neighbor from which the route was received. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|--|
| aggregate-address, on page 12 | Creates an aggregate entry in a BGP routing table. |
| bgp dampening, on page 52 | Enables BGP route dampening or changes various BGP route dampening factors. |
| clear bgp dampening, on page 85 | Clears BGP route dampening information and unsuppresses the suppressed routes. |

| Command | Description |
|---|---|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp flap-statistics, on page 242 | Displays BGP routes that have flapped. |
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |

show bgp flap-statistics

To display information about Border Gateway Protocol (BGP) paths that have flapped, use the **show bgp flap-statistics** command in XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] flap-statistics [{regex
regular-expression | route-policy route-policy-name | cidr-only | {ip-address | {mask /prefix-length}}}]
[longer-prefixes] [detail] [standby]
show bgp [ipv6 {unicast | labeled-unicast | all}] flap-statistics [{regex
regular-expression | route-policy route-policy-name | cidr-only | {ip-address | {mask /prefix-length}}}]
[longer-prefixes] [detail] [standby]
show bgp [all {unicast | labeled-unicast | all}] flap-statistics [{regex regular-expression | route-policy
route-policy-name | cidr-only | {ip-address | {mask /prefix-length}}}] [longer-prefixes] [detail]
[standby]
```

Syntax Description

| | |
|---|--|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| regex <i>regular-expression</i> | (Optional) Displays flap statistics for all paths that match the regular expression. |
| route-policy <i>route-policy-name</i> | (Optional) Displays flap statistics for a route policy. |
| cidr-only | (Optional) Displays only routes whose prefix length does not match the classful prefix length for that network. The cidr-only keyword can be specified only if the address family is IPv4. |
| <i>ip-address</i> | (Optional) Flap statistics for a network address only. |
| <i>mask</i> | (Optional) Network mask applied to the <i>ip-address</i> argument. |
| <i>/ prefix-length</i> | (Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| longer-prefixes | (Optional) Displays flap statistics for the specified prefix and more-specific prefixes. The longer-prefixes keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| detail | (Optional) Displays dampening parameters for the path. The detail keyword cannot be specified if the longer-prefixes keyword is specified. The detail keyword is available when the <i>ip-address</i> argument or <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |

| | |
|----------------|---|
| standby | (Optional) Displays information about the standby card. |
|----------------|---|

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Flap statistics are maintained only for paths if dampening is enabled using the **bgp dampening** command. If dampening is not enabled, the **show bgp flap-statistics** command does not display any paths.

If no arguments or keywords are specified, the software displays flap statistics for all paths for the specified address family. You can use the **regex**, **filter-list**, **cidr-only**, and **longer-prefixes** options to limit the set of paths displayed.

If you specify a network address without a mask or prefix length, the longest matching prefix for the network address is displayed. When displaying flap statistics for a single route, use the **detail** keyword to display dampening parameters for the route.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp flap-statistics** command:

```
RP/0/RP0/CPU0:router# show bgp flap-statistics
```

```
BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 26180
Dampening enabled
```

```

BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        From          Flaps Duration Reuse      Path
*d 10.0.0.0      172.20.16.177  4      00:13:31 00:18:10 100
*d 10.10.0.0     172.20.16.177  4      00:02:45 00:28:20 100

```

The following is sample output from the **show bgp flap-statistics** command with the **detail** keyword in XR EXEC mode:

```

RP/0/RP0/CPU0:router# show bgp flap-statistics 172.31.12.166 detail

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 738
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        From          Flaps Duration Reuse      Path
h 172.31.12.166      10.0.101.1    6      00:03:28          2 2000 3000

Half life      Suppress      Reuse penalty  Max. supp. time
00:15:00       2000          750            01:00:00

```

This table describes the significant fields shown in the display.

Table 12: show bgp flap-statistics Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP route identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network that is dampened. |
| From | IP address of the peer that advertised this route. |
| Flaps | Number of times the route has flapped. |
| Duration | Time (in hours:minutes:seconds) since the first flap. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path of the route that is being dampened. |
| Half life | Half-life value used when dampening this route. The half-life is the amount of time that must elapse to reduce the reuse penalty by half. The half-life value is specified using the bgp dampening command. |
| Suppress | Suppress value used to dampen this route. The suppress value is the value that the penalty must exceed for the route to be suppressed. The suppress value can be configured using the bgp dampening command. |

| Field | Description |
|----------------|--|
| Reuse penalty | Reuse penalty used to dampen this route. The penalty must fall below the reuse penalty for the route to be unsuppressed. The reuse penalty can be configured using the bgp dampening command. |
| Max supp. time | Maximum length of time that the route may be suppressed due to dampening. The maximum suppress time can be configured using the bgp dampening command. |

Related Commands

| Command | Description |
|--|---|
| bgp dampening, on page 52 | Enables BGP route dampening or changes various BGP route dampening factors. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp dampened-paths, on page 238 | Displays the BGP dampened routes. |
| show bgp neighbors, on page 256 | Displays information about BGP neighbors. |

show bgp inconsistent-as

To display Border Gateway Protocol (BGP) routes originated from more than one autonomous system, use the **show bgp inconsistent-as** command in XR EXEC mode.

```
show bgp [ipv4 {unicast | labeled-unicast | all}] inconsistent-as [standby]
show bgp [ipv6 {unicast | labeled-unicast | all}] inconsistent-as [standby]
show bgp [all {unicast | labeled-unicast | all}] inconsistent-as [standby]
```

| | | |
|---------------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

Use the **show bgp inconsistent-as** command to search through all prefixes in the specified BGP routing table and display the paths for any prefix that has inconsistent originating autonomous system numbers. The originating autonomous system is the last autonomous system number displayed in the path field and should be the same for all paths.

If a prefix has one or more paths originating from different autonomous systems, all paths for that prefix are displayed.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp inconsistent-as** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp inconsistent-as

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 1129
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop          Metric      LocPrf Weight Path
* 10.0.0.0      172.16.232.55          0           0 300 88 90 99 ?
*>             172.16.232.52      2222           0 400 ?
* 172.16.0.0    172.16.232.55          0           0 300 90 99 88 200 ?
*>             172.16.232.52      2222           0 400 ?
* 192.168.199.0 172.16.232.55          0           0 300 88 90 99 ?
*>             172.16.232.52      2222           0 400 ?
```

This table describes the significant fields shown in the display.

Table 13: show bgp inconsistent-as Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|---|
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default -safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |

show bgp labels

To display Border Gateway Protocol (BGP) routes and their incoming and outgoing labels, use the **show bgp labels** command in XR EXEC mode.

show bgp labels

| | | |
|---------------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled-unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read |
| Related Commands | Command | Description |
| | set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| | set default-safi | Sets the default subaddress Family Identifier (SAFI) for the current session. |

show bgp neighbor-group

To display information about the Border Gateway Protocol (BGP) configuration for neighbor groups, use the **show bgp neighbor-group** command in XR EXEC mode.

show bgp neighbor-group *group-name* {**configuration** [**defaults**] [**nvgen**] | **inheritance** | **users**}

Syntax Description

| | |
|----------------------|--|
| <i>group-name</i> | Name of the address family group to display. |
| configuration | (Optional) Displays the effective configuration for the neighbor group, including any configuration inherited by this neighbor group. |
| defaults | (Optional) Displays all configuration, including default configuration. |
| nvgen | (Optional) Displays output in show running-config command output. If the defaults keyword is also specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | Displays the af-groups, session groups, and neighbor groups from which this neighbor group inherits configuration. |
| users | Displays the neighbors and neighbor groups that inherit configuration from this neighbor group. |

Command Default

No default behavior or value

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp neighbor-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a neighbor group, including any configuration inherited from session groups, address family groups, and neighbor groups through application of the **use** command. The source of each configured command is also displayed.

Use the **defaults** keyword to display all configuration for the neighbor group, including default configuration. The command output identifies default onfiguration. Use the **nvgen** keyword to display configuration in the output form of **show running-config** command. Output in this form is suitable for cutting and pasting into a configuration session.

The **show bgp neighbor-group** command with the *group-name* **inheritance** argument and keyword displays the session groups, address family groups, and neighbor groups from which the specified neighbor group inherits configuration.

The **show bgp neighbor-group** *group-name* command displays the neighbors and neighbor groups that inherit configuration from the specified neighbor group.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The examples use the following configuration:

```
af-group group3 address-family ipv4 unicast
  remove-private-AS
  soft-reconfiguration inbound
!
af-group group2 address-family ipv4 unicast
  use af-group group3
  send-community-ebgp
  send-extended-community-ebgp
  capability orf prefix both
!
session-group group3
  dmzlink-bw
!
neighbor-group group3
  use session-group group3
  timers 30 90
!
neighbor-group group1
  remote-as 1982
  use neighbor-group group2
  address-family ipv4 unicast
!
!
neighbor-group group2
  use neighbor-group group3
  address-family ipv4 unicast
  use af-group group2
  weight 100
!
```

The following is sample output from the **show bgp neighbor-group** command with the **configuration** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group1 configuration
```

```
neighbor-group group1
  remote-as 1982          []
  timers 30 90           [n:group2 n:group3]
  dmzlink-bw             [n:group2 n:group3 s:group3]
  address-family ipv4 unicast []
  capability orf prefix both [n:group2 a:group2]
  remove-private-AS      [n:group2 a:group2 a:group3]
  send-community-ebgp    [n:group2 a:group2]
  send-extended-community-ebgp [n:group2 a:group2]
  soft-reconfiguration inbound [n:group2 a:group2 a:group3]
  weight 100             [n:group2]
```

The configuration source is shown to the right of each command. In the output, the **remote-as** command is configured directly on neighbor group group1, and the **send-community-ebgp** command is inherited from neighbor group group2, which in turn inherits the setting from af-group group2.

The following is sample output from the **show bgp neighbor-group** command with the **users** keyword. This output shows that the group1 neighbor group inherits session (address family-independent configuration parameters) from the group2 neighbor group. The group1 neighbor group also inherits IPv4 unicast configuration parameters from the group2 neighbor group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group2 users

Session:      n:group1
IPv4 Unicast: n:group1
```

The following is sample output from the **show bgp neighbor-group** command with the **inheritance** keyword. This output shows that the specified neighbor group group1 inherits session (address family-independent configuration) from neighbor group group2, which inherits its own session from neighbor group group3. Neighbor group group3 inherited its session from session group group3. It also shows that the group1 neighbor-group inherits IPv4 unicast configuration parameters from the group2 neighbor group, which in turn inherits them from the group2 af-group, which itself inherits them from the group3 af-group:

```
RP/0/RP0/CPU0:router# show bgp neighbor-group group1 inheritance

Session:      n:group2 n:group3 s:group3
IPv4 Unicast: n:group2 a:group2 a:group3
```

This table describes the significant fields shown in the display.

Table 14: show bgp neighbor-group Field Descriptions

| Field | Description |
|-----------|--|
| [] | Configures the command directly on the specified address family group. |
| s: | Indicates the name that follows is a session group. |
| a: | Indicates the name that follows is an address family group. |
| n: | Indicates the name that follows is a neighbor group. |
| [dflt] | Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the defaults keyword is specified. |
| <not set> | Indicates that the default is for the setting to be disabled. This field may be shown when the defaults keyword is specified. |

Related Commands

| Command | Description |
|--|--|
| af-group, on page 10 | Configures a BGP address family group. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

| Command | Description |
|---|--|
| show bgp af-group, on page 220 | Displays information about configuration for address family groups. |
| show bgp neighbors, on page 256 | Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups. |
| show bgp session-group, on page 316 | Displays information about the BGP configuration for session groups. |
| show running-config | Displays the contents of the currently running configuration or a subset of that configuration. |
| use, on page 368 | Inherits configuration from a neighbor group, a session group, or an address family group. |

show bgp neighbors

To display information about Border Gateway Protocol (BGP) connections to neighbors, use the **show bgp neighbors** command in XR EXEC mode.

```
show bgp neighbors [{performance-statistics | missing-eor}] [standby]
show bgp neighbors ip-address [{advertised-routes | dampened-routes | flap-statistics |
performance-statistics | received | {prefix-filter | routes} | routes}] [standby]
show bgp neighbors ip-address [{configuration | defaults} | nvgen | inheritance][standby]
show bgp neighbors ip-address decoded-message-log [{in | out}] [standby]
```

Syntax Description

| | |
|--|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| performance-statistics | (Optional) Displays performance statistics relative to work done by the BGP process for this neighbor. |
| missing-eor | (Optional) Displays neighbors that did not send end-of-rib (EoR) notification in read-only mode. |
| <i>ip-address</i> | (Optional) IP address of the BGP-speaking neighbor. If you omit this argument, all neighbors are displayed. |
| advertised-routes | (Optional) Displays all routes the router advertised to the neighbor. |
| dampened-routes | (Optional) Displays the dampened routes that are learned from the neighbor. |
| flap-statistics | (Optional) Displays flap statistics of the routes learned from the neighbor. |
| received { prefix-filter routes } | (Optional) Displays information received from the BGP neighbor. The options are: prefix-filter — Displays the prefix list filter. routes —Displays routes from the neighbor before inbound policy |
| routes | (Optional) Displays routes learned from the neighbor. |
| configuration | (Optional) Displays the effective configuration for the neighbor, including any settings that have been inherited from session groups, neighbor groups, or af-groups used by this neighbor. |
| defaults | (Optional) Displays all configuration settings, including any default settings. |

| | |
|----------------------------|--|
| nvgen | (Optional) Displays output in the show running-config command output. |
| inheritance | (Optional) Displays the session groups, neighbor groups, and af-groups from which this neighbor inherits configuration settings. |
| decoded-message-log | (Optional) Displays BGP message logs. |
| in | (Optional) Displays BGP inbound messages. |
| out | (Optional) Displays BGP outbound messages. |
| standby | (Optional) Displays information about the standby card. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 5.0.0 | This command was introduced. |
| Release 6.0 | The decoded-message-log [in out] option was added. |

Usage Guidelines**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify which routing table should be examined. If the **all** keyword is specified for address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp neighbors** command to display detailed information about all neighbors or a specific neighbor. Use the **performance-statistics** keyword to display information about the work related to specific neighbors done by the BGP process.

Use the **show bgp neighbors** command with the *ip-address* **received prefix-filter** argument and keyword to display the Outbound Route Filter (ORF) received from a neighbor.

Use the **advertised-routes** keyword to display a summary of the routes advertised to the specified neighbor.

Use the **dampened-routes** keyword to display routes received from the specified neighbor that have been suppressed due to dampening. For more details, see the **show bgp dampened-paths** command.

To display information about flapping routes received from a neighbor, use the **flap-statistics** keyword. For more details, see the **show bgp flap-statistics** command.

To display the routes received from a neighbor, use the **routes** keyword. For more details, see the **show bgp** command.

Use the **show bgp neighbor** command with the *ip-address* **configuration** argument and keyword to display the effective configuration of a neighbor, including configuration inherited from session groups, neighbor groups, or af-groups through application of the **use** command. Use the **defaults** keyword to display the value of all configurations for the neighbor, including default configuration. Use the **nvgen** keyword to display configuration output format of the **show running-config** command. Output in this format is suitable for cutting and pasting into a configuration session. Use the **show bgp neighbors** command with the *ip-address* **inheritance** argument and keyword to display the session groups, neighbor groups, and af-groups from which the specified neighbor inherits configuration.

Task ID

Task Operations ID

bgp read

The following is the sample output from the **show bgp neighbors** command with the *ip-address* and **configuration** argument and keyword to display graceful maintenance feature attributes:

```
*****
RP/0/0/CPU0:R1#show bgp neighbor 12.12.12.5
...
  Graceful Maintenance locally active, Local Pref=45, AS prepends=3
...
  For Address Family: IPv4 Unicast
...
  GSHUT Community attribute sent to this neighbor
...

*****

RP/0/0/CPU0:R1#show bgp neighbor 12.12.12.5 configuration
Mon Feb  2 14:30:41.042 PST
neighbor 12.12.12.5
  remote-as 1 []
  graceful-maintenance 1 []
  gr-maint local-preference 45 []
  gr-maint as-prepend 3 []
  gr-maint activate []

*****
```

Examples

The following is sample output from the **show bgp neighbors** command:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.1

BGP neighbor is 10.0.101.1, remote AS 2, local AS 1, external link
  Description: routem neighbor
  Remote router ID 10.0.101.1
  BGP state = Established, up for 00:00:56
  TCP open mode: passive only
BGP neighbor is 1.1.1.2
  Remote AS 300, local AS 100, external link
  Remote router ID 0.0.0.0
```

```

BGP state = Idle (LC/FIB for the neighbor in reloading)
Last read 00:00:00, Last read before reset 00:05:12
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3

BFD enabled (session initializing)
Last read 00:00:55, hold time is 180, keepalive interval is 60 seconds
DMZ-link bandwidth is 1000 Mb/s
Neighbor capabilities:
  Route refresh: advertised
  4-byte AS: advertised and received
  Address family IPv4 Unicast: advertised and received

Received 119 messages, 0 notifications, 0 in queue
Sent 119 messages, 22 notifications, 0 in queue
Minimum time between advertisement runs is 60 seconds

For Address Family: IPv4 Unicast
BGP neighbor version 137
Update group: 1.3
Community attribute sent to this neighbor
AF-dependant capabilities:
  Outbound Route Filter (ORF) type (128) Prefix-list:
    Send-mode: advertised
    Receive-mode: advertised
Route refresh request: received 0, sent 0
Policy for incoming advertisements is pass-all
Policy for outgoing advertisements is pass-all
5 accepted prefixes, 5 are bestpaths
Prefix advertised 3, suppressed 0, withdrawn 0, maximum limit 1000000
Threshold for warning message 75%

```

This table describes the significant fields shown in the display.

Table 15: show bgp neighbors Field Descriptions

| Field | Description |
|--------------|---|
| BGP neighbor | IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external. |
| Description | Neighbor specific description. |
| remote AS | <ul style="list-style-type: none"> • Number of the autonomous system to which the neighbor belongs. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. |

| Field | Description |
|--------------------------------------|---|
| local AS | Autonomous system number of the local system. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| internal link | Neighbor is an internal BGP peer. |
| external link | Neighbor is an external BGP peer. |
| Administratively shut down | Neighbor connection is disabled using the shutdown command. |
| remote router ID | Router ID (an IP address) of the neighbor. |
| Neighbor under common administration | Neighbor is internal or a confederation peer. |
| BGP state | Internal state of this BGP connection. |
| BFD enabled | Status of bidirectional forwarding detection. |
| TCP open mode | TCP mode used in establishing the BGP session. The following valid TCP mode are supported: <ul style="list-style-type: none"> • default—Accept active/passive connections • passive-only—Accept only passive connections • active-only—Accept only active connections initiated by the router |
| Last read | Time since BGP last read a message from this neighbor. |
| hold time | Hold time (in seconds) used on the connection with this neighbor. |
| keepalive interval | Interval for sending keepalives to this neighbor. |
| DMZ-link bandwidth | DMZ link bandwidth for this neighbor. |
| Neighbor capabilities | BGP capabilities advertised and received from this neighbor. The following valid BGP capabilities are supported: <ul style="list-style-type: none"> • Multi-protocol • Route refresh • Graceful restart • Outbound Route Filter (ORF) type (128) Prefix |
| Route refresh | Indicates that the neighbor supports dynamic soft reset using the route refresh capability. |
| 4-byte AS | Indicates that the neighbor supports the 4-byte AS capability. |

| Field | Description |
|--|---|
| Address family | Indicates that the local system supports the displayed address family capability. If “received” is displayed, the neighbor also supports the displayed address family. |
| Received | Number of messages received from this neighbor, the number of notification messages received and processed from this neighbor, and the number of messages that have been received, but not yet processed. |
| Sent | Number of messages sent to this neighbor, the number of notification messages generated to be sent to this neighbor, and the number of messages queued to be sent to this neighbor. |
| Minimum time between advertisement runs | Advertisement interval (in seconds) for this neighbor. |
| For Address Family | Information that follows is specific to the displayed address family. |
| BGP neighbor version | Last version of the BGP database that was sent to the neighbor for the specified address family. |
| Update group | Update group to which the neighbor belongs. |
| Route reflector client | Indicates that the local system is acting as a route reflector for this neighbor. |
| Inbound soft reconfiguration allowed | Indicates that soft reconfiguration is enabled for routes received from this neighbor. Note If the neighbor has route refresh capability, then soft configuration received-only routes are not stored by the local system unless “override route refresh” is displayed. |
| eBGP neighbor with no inbound or outbound policy: defaults to drop | Indicates that the neighbor does not have an inbound or outbound policy configured using the route-policy (BGP) command. Hence, no routes are accepted from or advertised to this neighbor. |
| Private AS number removed from updates to this neighbor | Indicates that remove-private-AS is configured on the specified address family for this neighbor. |
| NEXT_HOP is always this router | Indicates that next-hop-self is configured on the specified address family for this neighbor. |
| Community attribute sent to this neighbor | Indicates that send-community-ebgp is configured on the specified address family for this neighbor. |
| Extended community attribute sent to this neighbor | Indicates that send-extended-community-ebgp is configured on the specified address family for this neighbor. |
| Default information originate | Indicates that default-originate is configured on the specified address family for this neighbor, together with the policy used, if one was specified in the default-originate configuration. An indication of whether the default route has been advertised to the neighbor is also shown. |

| Field | Description |
|---|---|
| AF-dependant capabilities | <p>BGP capabilities that are specific to a particular address family. The following valid AF-dependent BGP capabilities are supported:</p> <ul style="list-style-type: none"> • route refresh capability • route refresh capability OLD value |
| Outbound Route Filter | <p>Neighbor has the Outbound Route Filter (ORF) capability for the specified address family. Details of the capabilities supported are also shown:</p> <p>Send-mode—“advertised” is shown if the local system can send an outbound route filter to the neighbor. “received” is shown if the neighbor can send an outbound route filter to the local system.</p> <p>Receive-mode—“advertised” is shown if the local system can receive an outbound route filter from the neighbor. “received” is shown if the neighbor can receive an outbound route filter from the local system.</p> |
| Graceful Restart Capability | Indicates whether graceful restart capability has been advertised to and received from the neighbor for the specified address family. |
| Neighbor preserved the forwarding state during latest restart | Indicates that when the neighbor connection was last established, the neighbor indicated that it preserved its forwarding state for the specified address family. |
| Local restart time | Restart time (in seconds) advertised to this neighbor. |
| RIB purge time | RIB purge time (in seconds) used for graceful restarts. |
| Maximum stalepath time | Maximum time (in seconds) a path received from this neighbor may be marked as stale if the neighbor restarts. |
| Remote Restart time | Restart time received from this neighbor. |
| Route refresh request | Number of route refresh requests sent and received from this neighbor. |
| Outbound Route Filter (ORF) | <p>“sent” indicates that an outbound route filter has been sent to this neighbor. “received” indicates that an outbound route filter has been received from this neighbor.</p> <p>Note A received outbound route filter may be displayed using the show bgp neighbors command with the received prefix-filter keywords.</p> |
| First update is deferred until ORF or ROUTE-REFRESH is received | If the local system advertised the receive capability and the neighbor has advertised send capability, no updates are generated until specifically asked by the neighbor (using a ROUTE-REFRESH or ORF with immediate request). |
| Scheduled to send the Prefix-list filter | Indicates the local system is due to send an outbound route filter request in order to receive updates from the neighbor. |
| Inbound path policy | Indicates if an inbound path policy is configured. |

| Field | Description |
|---|---|
| Outbound path policy | Indicates if an outbound path policy is configured. |
| Incoming update prefix filter list | Indicates a prefix list is configured to filter inbound updates from the neighbor. |
| Default weight | Default weight for routes received from the neighbor. |
| Policy for incoming advertisements | Indicates a route policy is configured to be applied to inbound updates from the neighbor. |
| Policy for outgoing advertisements | Indicates a route policy is configured to be applied to outbound updates to the neighbor. |
| Type | Indicates whether the condition map selects routes that should be advertised, or routes that should not be advertised: Exist—Routes advertised if permitted by the condition route map. Non-exist—Routes advertised if denied by the condition route map. |
| accepted prefixes | Number of prefixes accepted. |
| Prefix advertised | Number of prefixes advertised to the neighbor during the lifetime of the current connection with the neighbor. |
| suppressed | Number of prefix updates that were suppressed because no transitive attributes changed from one best path to the next. Note Update suppression occurs only for external BGP neighbors. |
| withdrawn | Number of prefixes withdrawn from the neighbor during the lifetime of the current connection with the neighbor. |
| maximum limit | Maximum number of prefixes that may be received from the neighbor. If “(warning-only)” is displayed, a warning message is generated when the limit is exceeded, otherwise the neighbor connection is shut down when the limit is exceeded. |
| Threshold for warning message | Percentage of maximum prefix limit for the neighbor at which a warning message is generated. |
| Connections established | Number of times the router has established a BGP peering session with the neighbor. |
| dropped | Number of times that a good connection has failed or been taken down. |
| Last reset due to | Reason that the connection with the neighbor was last reset. |
| Time since last notification sent to neighbor | Amount of time since a notification message was last sent to the neighbor. |
| Error Code | Type of notification that was sent. The notification data, if any, is also displayed. |

| Field | Description |
|---|---|
| Time since last notification received from neighbor | Amount of time since a notification message was last received from the neighbor. |
| Error Code | Type of notification that was received. The notification data received, if any, is also displayed |
| External BGP neighbor may be up to <n> hops away | Indicates ebgp-multihop is configured for the neighbor. |
| External BGP neighbor not directly connected | Indicates that the neighbor is not directly attached to the local system. |
| Notification data sent: | Data providing more details on the error along with the error notification sent to the neighbor. |

The following is sample output from the **show bgp neighbors** command with the **advertised-routes** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbors 172.20.16.178 routes

BGP router identifier 172.20.16.181, local AS number 1
BGP main routing table version 27
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*> 10.0.0.0       172.20.16.178          40             0 10 ?
*> 10.22.0.0      172.20.16.178          40             0 10 ?
```

The following is sample output from the **show bgp neighbors** command with the **routes** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.1 dampened-routes

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          From          Reuse      Path
*d 10.0.0.0       10.0.101.1    00:59:30 2 100 1000 i
*d 11.0.0.0       10.0.101.1    00:59:30 2 100 1000 i
*d 12.0.0.0       10.0.101.1    00:59:30 2 100 1000 i
*d 13.0.0.0       10.0.101.1    00:59:30 2 100 1000 i
*d 14.0.0.0       10.0.101.1    00:59:30 2 100 1000 i
```


This table describes the significant fields shown in the display.

Table 16: show bgp neighbors routes Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |

| Field | Description |
|----------|--|
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

The following is sample output from the **show bgp neighbors** command with the **dampened-routes** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.1 flap-statistics

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          From          Flaps  Duration  Reuse      Path
   ----          -
h 10.1.0.0          10.0.101.1      5008   2d02h           2 5000 1000
h 10.2.0.0          10.0.101.1      5008   2d02h           2 2000 3000
h 10.2.0.0          10.0.101.1      5008   2d02h           2 9000 6000
*d 10.0.0.0          10.0.101.1      5008   2d02h    00:59:30 2 100 1000
h 10.0.0.0/16        10.0.101.1      5008   2d02h           2 100 102
*d 10.11.0.0         10.0.101.1      5008   2d02h    00:59:30 2 100 1000
*d 10.12.0.0         10.0.101.1      5008   2d02h    00:59:30 2 100 1000
*d 10.13.0.0         10.0.101.1      5008   2d02h    00:59:30 2 100 1000
*d 10.14.0.0         10.0.101.1      5008   2d02h    00:59:30 2 100 1000
h 192.168.0.0/16     10.0.101.1      5008   2d02h           2 100 101
```

This table describes the significant fields shown in the display.

Table 17: show bgp neighbors dampened-routes Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |

| Field | Description |
|-------------------|--|
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| From | Neighbor from which the route was received. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

The following is sample output from the **show bgp neighbors** command with the **flap-statistics** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.2 performance-statistics

BGP neighbor is 10.0.101.2, remote AS 1
  Read 3023 messages (58639 bytes) in 3019 calls (time spent: 1.312 secs)
  Read throttled 0 times
```

```

Processed 3023 inbound messages (time spent: 0.198 secs)
Wrote 58410 bytes in 6062 calls (time spent: 3.041 secs)
Processing write list: wrote 0 messages in 0 calls (time spent: 0.000 secs)
Processing write queue: wrote 3040 messages in 3040 calls (time spent: 0.055 secs)

```

```

Received 3023 messages, 0 notifications, 0 in queue
Sent 3040 messages, 0 notifications, 0 in queue

```

This table describes the significant fields shown in the display.

Table 18: show bgp neighbors flap-statistics Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP route identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between when the BGP process scans for the specified address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |

| Field | Description |
|--------------|---|
| Origin codes | Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |
| Network | IP prefix and prefix length for a network. |
| From | IP address of the peer that advertised this route. |
| Flaps | Number of times the route has flapped. |
| Duration | Time (in hours:minutes:seconds) since the router noticed the first flap. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path to reach the destination network. |

The following is sample output from the **show bgp neighbors** command with the **performance-statistics** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.2 performance-statistics
BGP neighbor is 10.0.101.2, remote AS 1
  Read 3023 messages (58639 bytes) in 3019 calls (time spent: 1.312 secs)
  Read throttled 0 times
  Processed 3023 inbound messages (time spent: 0.198 secs)
  Wrote 58410 bytes in 6062 calls (time spent: 3.041 secs)
  Processing write list: wrote 0 messages in 0 calls (time spent: 0.000 secs)
  Processing write queue: wrote 3040 messages in 3040 calls (time spent: 0.055 secs)
  Received 3023 messages, 0 notifications, 0 in queue
  Sent 3040 messages, 0 notifications, 0 in queue
```

This table describes the significant fields shown in the display.

Table 19: show bgp neighbors performance-statistics Field Descriptions

| Field | Description |
|------------------|--|
| Read | Indicates the number of messages received from the neighbor, the total size of received messages, the number of read operations performed, and the real time spent (in seconds) by the process performing read operations for this neighbor. |
| Read throttled | Number of times that reading from the TCP connection to this neighbor has been throttled. Throttling is due to a backlog of messages that have been read but not processed. |
| inbound messages | Number of read messages that have been processed, and the real time spent processing inbound messages for this neighbor. |

| Field | Description |
|------------------------|---|
| Wrote | Amount of data that has been sent to this neighbor, number of write operations performed, and the real time spent by the process performing write operations for this neighbor. |
| Processing write list | Number of messages written from the write list to this neighbor, number of times the write list has been processed, and real time spent processing the write list. Note Write lists typically contain only update messages. |
| Processing write queue | Number of messages written from the write queue to this neighbor, number of times the write queue has been processed, and real time spent processing the write queue. |
| Received | Number of messages received from this neighbor, number of notification messages received and processed from this neighbor, and number of messages that have been received, but not yet processed. |
| Sent | Number of messages sent to this neighbor, number of notification messages generated to be sent to this neighbor, and number of messages queued to be sent to this neighbor. |

The following is sample output from the **show bgp neighbors** command with the **configuration** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbors 10.0.101.1 configuration

neighbor 10.0.101.1
  remote-as 2                []
  bfd fast-detect            []
  address-family ipv4 unicast []
  policy pass-all in         []
  policy pass-all out        []

  policy pass-all in         []
  policy pass-all out        []
```

This table describes the significant fields shown in the display.

Table 20: show bgp neighbors configuration Field Descriptions

| Field | Description |
|---------------------------|--|
| neighbor | IP address configuration of the neighbor. |
| remote-as | Remote autonomous system configured on the neighbor. |
| bfd fast-detect | BFD parameter configured on the neighbor. |
| address-family | Address family and subsequent address family configured on the router. |
| route-policy pass-all in | Route policy configured for inbound updates. |
| route-policy pass-all out | Route policy configured for outbound updates. |

This sample output from the **show bgp neighbor** command displays TCP MSS information for the specified neighbor:

```
RP/0/RP0/CPU0:router#show bgp neighbor 10.0.0.2

BGP neighbor is 10.0.0.2
Remote AS 1, local AS 1, internal link
Remote router ID 10.0.0.2
BGP state = Established, up for 00:09:17
Last read 00:00:16, Last read before reset 00:00:00
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:16, attempted 19, written 19
Second last write 00:01:16, attempted 19, written 19
Last write before reset 00:00:00, attempted 0, written 0
Second last write before reset 00:00:00, attempted 0, written 0
Last write pulse rcvd Dec 7 11:58:42.411 last full not set pulse count 23
Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 00:00:00, second last 00:00:00
Precedence: internet
Multi-protocol capability received
Neighbor capabilities:
Route refresh: advertised (old + new) and received (old + new)
Graceful Restart (GR Awareness): advertised and received
4-byte AS: advertised and received
Address family IPv4 Unicast: advertised and received
Received 12 messages, 0 notifications, 0 in queue
Sent 12 messages, 0 notifications, 0 in queue
Minimum time between advertisement runs is 0 secs
TCP Maximum Segment Size 500

For Address Family: IPv4 Unicast
BGP neighbor version 4
Update group: 0.2 Filter-group: 0.1 No Refresh request being processed
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was received during read-only mode
Last ack version 4, Last synced ack version 0
Outstanding version objects: current 0, max 0
Additional-paths operation: None
Send Multicast Attributes
```

This sample output from the **show bgp neighbor** command with the **configuration** keyword displays TCP MSS configuration:

```
RP/0/RP0/CPU0:router#show bgp neighbor 10.0.0.2 configuration

neighbor 10.0.0.2
remote-as 1 []
tcp-mss 400 [n:n1]
```

```
address-family IPv4 Unicast []
```

Related Commands

| Command | Description |
|--|---|
| clear bgp , on page 82 | Resets a BGP connection or session. |
| network (BGP) , on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp , on page 193 | Displays entries in the BGP routing table. |
| show bgp dampened-paths , on page 238 | Displays BGP dampened routes. |
| show bgp flap-statistics , on page 242 | Displays BGP routes that have flapped. |
| show bgp neighbor-group , on page 252 | Displays information about the BGP configuration for neighbor groups. |
| shutdown (BGP) , on page 339 | Disables a neighbor without removing all of its configuration. |

show bgp neighbors nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information across neighbors, use the **show bgp neighbors nsr** command in XR EXEC mode.

show bgp [{**ipv4** {**unicast** | **all**} | **ipv6** {**unicast** | **all**} | **unicast** | **unicast**}] **neighbors nsr** [**standby**]

| | | |
|--------------------|----------------|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | standby | (Optional) Displays information about the standby card. |

| | |
|-----------------|--------------------------------|
| Command Default | No default behavior or values. |
|-----------------|--------------------------------|

| | |
|---------------|---------|
| Command Modes | XR EXEC |
|---------------|---------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.2.1 | This command was introduced. |

| | |
|------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|------------------|---|

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp neighbors nsr** command with the **standby** keyword:

```
RP/0/RP0/CPU0:router# show bgp neighbors nsr standby
```

```
BGP neighbor is 2.2.2.2
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0
```

```
BGP neighbor is 10.0.101.5
  BGP state = Established, up for 05:19:00
  NSR state = NSR Ready
  Outstanding Postits: 0
```

```

BGP neighbor is 10.1.0.5
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0

```

This table describes the significant fields shown in the display.

Table 21: show bgp neighbors nsr Field Descriptions

| Field | Description |
|---------------------|---|
| BGP state | Displays BGP neighbor peering state. |
| NSR state | Displays BGP neighbor NSR state. |
| Outstanding Postits | Displays the postit counters of pending events. |

Related Commands

| Command | Description |
|---|---|
| nsr (BGP), on page 155 | Activates the Border Gateway Protocol (BGP) nonstop routing (NSR). |
| show bgp summary nsr, on page 324 | Displays the Border Gateway Protocol (BGP) nonstop routing (NSR) information. |
| show bgp summary, on page 320 | Displays the status of all Border Gateway Protocol (BGP) connections. |

show bgp nexthops

To display statistical information about the Border Gateway Protocol (BGP) next hops, use the **show bgp nexthops** command in XR EXEC mode.

show bgp nexthops **A.B.C.D.aigp-value**[**statistics**] [**speaker** *speaker-id*] [**standby**]

| Syntax Description | <div> <div>ipv4</div> <div>(Optional) Specifies IP Version 4 address prefixes.</div> </div> <div> <div>unicast</div> <div>(Optional) Specifies unicast address prefixes.</div> </div> <div> <div>labeled-unicast</div> <div>(Optional) Specifies labeled-unicast address prefixes.</div> </div> <div> <div>all</div> <div>(Optional) For subaddress families, specifies prefixes for all subaddress families.</div> </div> <div> <div>ipv6</div> <div>(Optional) Specifies IP Version 6 address prefixes.</div> </div> <div> <div>all</div> <div>(Optional) For address family, specifies prefixes for all address families.</div> </div> <div> <div>statistics</div> <div>(Optional) Specifies nexthop statistics.</div> </div> <div> <div>speaker <i>speaker-id</i></div> <div>(Optional) Specifies a speaker process ID.</div> </div> <div> <div>standby</div> <div>(Optional) Displays information about the standby card.</div> </div> | |
|--------------------|---|------------------------------|
| Command Default | No default behavior or value | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>The show bgp nexthops command displays statistical information about next-hop notifications, the time spent processing the notifications, and details about each next-hop that has been registered with the Routing Information Base (RIB).</p> <p>The next-hop information is displayed for all active speaker processes in distributed mode. Each speaker displays a set of next-hops that belongs to the prefixes received by the speaker and next hops that belong to best paths that were received by other speaker processes. Use the speaker <i>speaker-id</i> keyword and argument to display information for only the specified speaker process.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read |

Related Commands

| Command | Description |
|---|---|
| bgp redistribute-internal, on page 70 | Specifies the delay for triggering BGP next-hop calculations. |

show bgp nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information, use the **show bgp nsr** command in XR EXEC mode.

show bgp [{**ipv4** {**unicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **all** | **labeled-unicast**} | **all** {**unicast** | **all** | **labeled-unicast**} | **unicast**}] **nsr** [**standby**]

| | | |
|---------------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | standby | Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.2.1 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read |

Examples

The following is sample output from the **show bgp nsr** command:

```
RP/0/RP0/CPU0:router# show bgp nsr

Fri Jan 30 10:18:48.171 PST PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
```

```

Default Cluster ID: 10.1.0.1
Active Cluster IDs: 10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

```

```

Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGP notified
RIB has converged: version 1

```

===== Post Failover Summary for Active instance =====

| Node | Process | Read | Write | Inbound |
|--------------|------------------------|--------|-------------------|---------|
| node0_0_CPU0 | Speaker | 146.75 | 18.90 | 3.46 |
| Entered mode | Standby Ready | | : Jan 30 10:00:39 | |
| Entered mode | TCP NSR Setup | | : Jan 30 10:00:39 | |
| Entered mode | TCP NSR Setup Done | | : Jan 30 10:00:39 | |
| Entered mode | TCP Initial Sync | | : Jan 30 10:00:39 | |
| Entered mode | TCP Initial Sync Done | | : Jan 30 10:00:44 | |
| Entered mode | FPBSN processing done | | : Jan 30 10:00:44 | |
| Entered mode | Update processing done | | : Jan 30 10:00:44 | |
| Entered mode | BGP Initial Sync | | : Jan 30 10:00:44 | |
| Entered mode | BGP Initial Sync done | | : Jan 30 10:00:44 | |
| Entered mode | NSR Ready | | : Jan 30 10:00:44 | |

```

Current BGP NSR state - NSR Ready achieved at: Jan 30 10:00:44
NSR State READY notified to Redcon at: Jan 30 10:16:58

```

NSR Post Failover Summary:

QAD Statistics:

| | | | |
|--------------------|-------|------------------------|-------|
| Messages Sent | : 512 | ACKs Received | : 512 |
| Messages Received | : 8 | ACKs Sent | : 8 |
| Send Failures | : 1 | Send ACK Failures | : 0 |
| Suspends | : 1 | Resumes | : 1 |
| Messages Processed | : 8 | Out of sequence drops: | 0 |

Postit Summary:

```

Total pending postit messages: 0
Neighbors with pending postits: 0

```

```
Rib Trigger: enabled
Last RIB down event Jan 29 09:50:03.069 received
Last RIB convergence Jan 29 09:50:03.069 last ack received.
```

```
Address Family IPv4 Unicast converged in 87531 seconds
```

The following example shows sample output from the **show bgp nsr** command with the **standby** keyword:

```
RP/0/RP0/CPU0:router# show bgp nsr standby
```

```
Fri Jan 30 10:18:55.654 PST PDT
```

```
BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
Default Cluster ID: 10.1.0.1
Active Cluster IDs: 10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60
```

```
Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGP notified
RIB has converged: version 1
```

```
===== Post Failover Summary for Standby instance =====
```

| Node | Process | Read | Write | Inbound |
|--------------|--------------------|------|-----------------|---------|
| node0_1_CPU0 | Speaker | 1.68 | 0.00 | 1.42 |
| Entered mode | Standby Ready | : | Jan 30 10:00:39 | |
| Entered mode | TCP Replication | : | Jan 30 10:00:39 | |
| Entered mode | TCP Init Sync Done | : | Jan 30 10:00:44 | |
| Entered mode | NSR Ready | : | Jan 30 10:00:44 | |

```
QAD Statistics:
```

| | | | | |
|--------------------|-------|--------------------|-------|-----------------|
| Messages Sent | : 9 | ACKs Received | : 9 | |
| Messages Received | : 512 | ACKs Sent | : 512 | |
| Send Failures | : 0 | Send ACK Failures | : 0 | |
| Suspends | : 0 | Resumes | : 0 | |
| Messages Processed | : 512 | Standby init drops | : 0 | Out of sequence |

drops: 0

Postit Summary:

Total pending postit messages: 0
Neighbors with pending postits: 0

Rib Trigger: enabled

Last RIB down event Jan 29 09:50:17.308 received

Last RIB convergence Jan 29 09:50:17.308 last ack received.

Related Commands

| Command | Description |
|--|--|
| nsr (BGP), on page 155 | Activates Border Gateway Protocol (BGP) nonstop routing (NSR). |

show bgp paths

To display all the Border Gateway Protocol (BGP) paths in the database, use the **show bgp paths** command in XR EXEC mode.

show bgp paths [**detail**] [**debug**] [**regex** *regular-expression*]

| | | |
|--------------------|--|---|
| Syntax Description | detail | (Optional) Displays detailed attribute information. |
| | debug | (Optional) Displays attribute process ID, hash bucket, and hash chain ID attribute information. |
| | regex <i>regular-expression</i> | (Optional) Specifies an autonomous system path that matches the regular expression. |

Command Default No default behavior or values

Command Modes XR EXEC

| | | |
|-----------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp paths** command to display information about AS paths and the associated attributes with which the paths were received.

If no options are specified, all stored AS paths are displayed with the number of routes using each path.



Note The AS path information is stored independently of the address family, making it possible that routes from different address families could be using the same path.

Use the *regular-expression* argument to limit the output to only those paths that match the specified regular expression. See the for information on regular expressions.

Use the **detail** keyword to display detailed information on the attributes stored with the AS path.

| | | |
|---------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read |

Examples The following is sample output from the **show bgp paths** command:

```
RP/0/RP0/CPU0:router# show bgp paths detail
```

| Proc | Attributes | Refcount | Metric | Path |
|-------|---------------------------|----------|--------|-----------|
| Spk 0 | ORG AS LOCAL | 7 | 0 | i |
| Spk 0 | ORG AS LOCAL COMM EXTCOMM | 3 | 0 | 21 i |
| Spk 0 | MET ORG AS | 3 | 55 | 2 i |
| Spk 0 | ORG AS | 3 | 0 | 2 10 11 i |
| Spk 0 | ORG AS COMM | 3 | 0 | 2 10 11 i |
| Spk 0 | MET ORG AS ATOM | 3 | 2 | 2 3 4 ? |
| Spk 0 | MET ORG AS | 3 | 1 | 2 3 4 e |
| Spk 0 | MET ORG AS | 3 | 0 | 2 3 4 i |

This table describes the significant fields shown in the display.

Table 22: show bgp paths Field Descriptions

| Field | Description |
|------------|---|
| Proc | ID of the process in which the path is stored. This is always “Spk 0.” |
| Attributes | Attributes that are present. The following may appear: MET—Multi Exit Discriminator (MED) attribute is present. ORG—Origin attribute is present. AS—AS path attribute is present. LOCAL—Local preference attribute is present. AGG—Aggregator attribute is present. COMM—Communities attribute is present. ATOM—Atomic aggregate attribute is present. EXTCOMM—Extended communities attribute is present. |
| NeighborAS | Autonomous system number of the neighbor, or 0, if the path information originated locally. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| Refcount | Number of routes using a path. |
| Metric | Value of the interautonomous system metric, otherwise known as the MED metric. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |

show bgp policy

To display information about Border Gateway Protocol (BGP) advertisements under a proposed policy, use the **show bgp policy** command in XR EXEC mode.

show bgp policy

| | | |
|---------------------------|--|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | <i>ip-address</i> | (Optional) IP address of a single neighbor. |
| | sent-advertisements | (Optional) Displays the routes that have been advertised to neighbors. If a route has not yet been advertised to the neighbor, it is not shown. |
| | route-policy | (Optional) Displays advertisements for an output route policy. |
| | route-policy-name | (Optional) Name of the route policy. |
| | standby | (Optional) Displays information about the standby card. |
| | summary | (Optional) Displays a summary of the BGP advertisements. |
| Command Default | Advertisements for all neighbors are displayed if the neighbor <i>ip-address</i> keyword and argument are not specified. If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp policy** command to display routes that would be advertised to neighbors under a proposed policy. Unlike in the **show bgp advertised** command, the information displayed reflects any modifications made to the routes when executing the specified policy.

Use the **neighbor** keyword to limit the output to routes advertised to a particular neighbor. Use the **sent-advertisements** keyword to change the output in two ways:

- If a policy is not specified explicitly, any policy configured on the neighbor (using the **route-policy (BGP)** command) is executed before displaying the routes.
- Only routes that have already been advertised to the neighbor (and not withdrawn) are displayed. Routes that have not yet been advertised are not displayed.

Use the **summary** keyword to display abbreviated output.

Task ID

Task ID Operations

bgp read

Examples

The following is sample output from the **show bgp policy** command with the **summary** keyword in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp policy summary
```

| Network | Next Hop | From | Advertised to |
|---------------|------------|------------|--|
| 172.16.1.0/24 | 10.0.101.1 | 10.0.101.1 | 10.0.101.2 10.0.101.3 |
| 172.17.0.0/16 | 0.0.0.0 | Local | 10.0.101.1 10.0.101.2 10.0.101.3 |

This table describes the significant fields shown in the display.

Table 23: show bgp policy summary Field Descriptions

| Field | Description |
|---------|--|
| Network | IP prefix and prefix length for a network. |

| Field | Description |
|-----------------|--|
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| From | IP address of the peer that advertised this route. |
| Local | Indicates the route originated on the local system. |
| Local Aggregate | Indicates the route is an aggregate created on the local system. |
| Advertised to | Indicates the neighbors to which this route was advertised. |

The following is sample output from the **show bgp policy** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp policy

11.0.0.0/24 is advertised to 10.4.101.1
  Path info:
    neighbor: Local          neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath: 1

11.0.0.0/24 is advertised to 10.4.101.2
  Path info:
    neighbor: Local          neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath:

11.0.0.0/24 is advertised to 10.4.101.3
  Path info:
    neighbor: Local          neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
```

```

    aspath:

12.0.0.0/24 is advertised to 10.4.101.2
Path info:
  neighbor: 10.4.101.1      neighbor router id: 10.4.101.1
  valid external best
Attributes after inbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath: 2 3 4
Attributes after outbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath:2 3 4

12.0.0.0/24 is advertised to 10.4.101.3
Path info:
  neighbor: 10.4.101.1      neighbor router id: 10.4.101.1
  valid external best
Attributes after inbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath: 2 3 4
Attributes after outbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath:2 3 4

```

This table describes the significant fields shown in the display.

Table 24: show bgp policy Field Descriptions

| Field | Description |
|-------------------------------------|---|
| Is advertised to | IP address of the peer to which this route is advertised. If the route is advertised to multiple peers, information is shown separately for each peer. |
| neighbor | IP address of the peer that advertised this route, or one of the following: Local—Route originated on the local system. Local Aggregate—Route is an aggregate created on the local system. |
| neighbor router id | BGP identifier for the peer, or the local system if the route originated on the local system. |
| Not advertised to any peer | Indicates the no-advertise well-known community is associated with this route. Routes with this community are not advertised to any BGP peers. |
| Not advertised to any EBGp peer | Indicates the no-export well-known community is associated with this route. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router. |
| Not advertised outside the local AS | Indicates the local-AS well-known community is associated with this route. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. |

| Field | Description |
|-----------------------------|--|
| (Received from a RR-client) | Path was received from a route reflector client. |
| (received-only) | Path is not used for routing purposes. It is used to support soft reconfiguration, and records the path attributes before inbound policy was applied to a path received from a peer. A path marked “received-only” indicates that either the path was dropped by inbound policy, or that a copy of path information was created and then modified for routing use. |
| (received & used) | Indicates that the path is used both for soft reconfiguration and routing purposes. A path marked “(received & used)”, implies the path information was not modified by inbound policy. |
| valid | Path is valid. |
| redistributed | Path is locally sourced through redistribution. |
| aggregated | Path is locally sourced through aggregation. |
| local | Path is locally sourced through the network command. |
| confed | Path was received from a confederation peer. |
| best | Path is selected as best. |
| multipath | Path is one of multiple paths selected for load-sharing purposes. |
| dampinfo | Indicates dampening information: Penalty—Current penalty for this path. Flapped—Number of times the route has flapped. In—Time (hours:minutes:seconds) since the network first flapped. Reuse in—Time (hours:minutes:seconds) after which the path is available. This field is displayed only if the path is currently suppressed. |

| Field | Description |
|---|---|
| Attributes after inbound policy was applied | <p>Displays attributes associated with the received route, after any inbound policy has been applied.</p> <p>AGG—Aggregator attribute is present.</p> <p>AS—AS path attribute is present.</p> <p>ATOM—Atomic aggregate attribute is present.</p> <p>COMM—Communities attribute is present.</p> <p>EXTCOMM—Extended communities attribute is present.</p> <p>LOCAL—Local preference attribute is present.</p> <p>MET—Multi Exit Discriminator (MED) attribute is present.</p> <p>next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.</p> <p>ORG—Origin attribute is present.</p> |
| origin | <p>Origin of the path:</p> <p>IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>EGP—Path originated from an Exterior Gateway Protocol.</p> <p>incomplete—Origin of the path is not clear; in example, a route that is redistributed into BGP from an IGP.</p> |
| neighbor as | First autonomous system (AS) number in the AS path. |
| aggregator | Indicates that the path was received with the aggregator attribute. The AS number and router-id of the system that performed the aggregation are shown. |
| metric | Value of the interautonomous system metric, otherwise known as the MED metric. |
| localpref | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| aspath | AS path associated with the route. |
| community | <p>Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities:</p> <p>Local-AS—Community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise—Community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export—Community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.</p> |

| Field | Description |
|--|---|
| Extended community | Extended community attributes associated with the path. For known extended community types, the following codes may be displayed: RT—Route target community SoO—Site of Origin community LB—Link Bandwidth community |
| Originator | Router ID of the originating router when route reflection is used. |
| Cluster lists | Router ID or cluster ID of all route reflectors through which the route has passed. |
| Attributes after outbound policy was applied | Displays attributes associated with the received route, after any outbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present. |


Related Commands

| Command | Description |
|--|---|
| route-policy (BGP), on page 178 | Applies an inbound or outbound routing policy to a neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp advertised, on page 214 | Displays routes advertised to neighbors. |
| show bgp neighbors, on page 256 | Displays information about the TCP and BGP connections to neighbors. |
| show bgp route-policy, on page 312 | Displays BGP information about networks that match an outbound route policy. |

show bgp process

To display Border Gateway Protocol (BGP) process information, use the **show bgp process** command in XR EXEC mode.

show bgp process [**performance-statistics**] [**detail**] [**standby**]

| | | |
|--|--|--|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4. |
| | unicast | (Optional) Specifies the unicast subaddress family. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | performance- statistics | (Optional) Displays performance statistics relative to the work done by the specified process. |
| | detail | (Optional) Specifies detailed process information. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
|  | | |
| Note | The set default-afi command is used to specify the default address family for the session, and the set default-safi command is used to specify the default subaddress family for the session. See the <i>System Management Command Reference for Cisco NCS 6000 Series Routers</i> for detailed information and syntax for the set default-afi and set default-safi commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast. | |
| Use the show bgp process command to display status and summary information for the Border Gateway Protocol (BGP) process. The output shows various global and address family-specific BGP configurations. | | |

A summary of the number of neighbors, update messages, and notification messages sent and received by the process is also displayed.

Use the **detail** keyword to display detailed process information. The detailed process information shows the memory used by each of various internal structure types.

Use the **performance-statistics** keyword to display a summary or detail of work done by the BGP processes. The summary display shows the real time spent performing certain operations and the time stamps for state transitions during initial convergence.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp process** command:

```
RP/0/RP0/CPU0:router# show bgp process
```

```
BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60
```

```
Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Scan interval: 60
Main Table Version: 150
IGP notification: IGP notified
```

```
Node          Process      Nbrs Estab Rst Upd-Rcvd Upd-Sent Nfn-Rcvd Nfn-Sent
node0_0_CPU0 Speaker      3      2   1      20      10        0        0
```

This table describes the significant fields shown in the display.

Table 25: show bgp process Field Descriptions

| Field | Description |
|---------------------|---|
| BGP is operating in | Indicates BGP is operating in standalone mode. This is the only supported mode. |

| Field | Description |
|--|--|
| Autonomous System | Autonomous system number for the local system. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If no global ID is available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | Cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. This is controlled by the default-metric command. |
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. This is controlled by the bgp fast-external-fallover disable command. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the bgp log neighbor changes disable command. |
| Enforce first AS enabled | Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled. This is controlled by the bgp enforce-first-as disable command. |
| iBGP to IGP redistribution | Indicates internal redistribution is enabled using the bgp redistribution-internal command. |
| Treating missing MED as worst | Indicates missing Multi Exit Discriminator (MED) metric values are treated as worst in the route selection algorithm. This is controlled by the bgp bestpath med missing-as-worst command. |
| Always compare MED is enabled | Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the bgp bestpath med always command. |
| AS Path ignore is enabled | Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the bgp bestpath as-path ignore command. |
| Comparing MED from confederation peers | Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the bgp bestpath med confed command. |

| Field | Description |
|------------------------------------|--|
| Comparing router ID for eBGP paths | Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the bgp bestpath compare-routerid command. |
| Default local preference | Default local preference value used for BGP routes. This is controlled by the bgp default local-preference command. |
| Default keepalive | Default keepalive interval. This is controlled by the timers bgp command. |
| Graceful restart enabled | Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: <ul style="list-style-type: none"> • bgp graceful-restart • bgp graceful-restart purge-time • bgp graceful-restart stalepath-time • bgp graceful-restart restart-time • bgp graceful-restart graceful-reset |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the bgp scan-time command. |
| Dampening | Indicates whether dampening is enabled for the specified address family. This is controlled by the dampening command. |
| Client reflection | Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the bgp client-to-client reflection disable command. |
| Scan interval | Interval (in seconds) between BGP scans for the given address family. This is controlled by the bgp scan-time command in address family configuration mode. |
| Main Table Version | Last version of the BGP database that was installed into the main routing table. |
| IGP notification | Indicates whether Interior Gateway Protocols (IGP) have been notified of BGP convergence for the specified address family. |
| Node | Node on which the process is executing. |
| Process | Type of BGP process. |
| Speaker | Speaker process. A speaker process is responsible for receiving, processing, and sending BGP messages to configured neighbors. |
| Nbrs | Number of neighbors for which the process is responsible. |
| Estab | Number of neighbors that have connections in the established state for this process. |

| Field | Description |
|----------|--|
| Rst | Number of times this process was restarted. |
| Upd-Rcvd | Number of update messages received by the process. |
| Upd-Sent | Number of update messages sent by the process. |
| Nfn-Rcvd | Number of notification messages received by the process. |
| Nfn-Sent | Number of notification messages sent by the process. |

The following is sample output from the **show bgp process** command with the Graceful Maintenance feature enabled:

```
RP/0/0/CPU0:R1#show bgp process
```

```
...
```

```
Graceful Maintenance active. Retaining routes in RIB during BGP shutdown
```

```
...
```

Or

```
Graceful Maintenance active for all neighbors. Retaining routes in RIB during BGP shutdown
```

```
*****
```

```
RP/0/0/CPU0:Jan 28 22:01:36.356 : bgp[1056]: %ROUTING-BGP-5-ADJCHANGE : neighbor 10.10.10.4
Up (VRF: default) (AS: 4) WARNING: Graceful Maintenance is Active
```

The following is sample output from the **show bgp process** command with the **detail** keyword:

```
RP/0/RP0/CPU0:router# show bgp all all process detail
```

```
BGP Process Information
```

```
BGP is operating in STANDALONE mode
```

```
Autonomous System: 1
```

```
Router ID: 10.0.0.5 (manually configured)
```

```
Cluster ID: 10.0.0.5
```

```
Fast external fallover enabled
```

```
Neighbor logging is enabled
```

```
Enforce first AS enabled
```

```
Default local preference: 100
```

```
Default keepalive: 60
```

```
Update delay: 120
```

```
Generic scan interval: 60
```

```
BGP Speaker process: 0, location node0_0_0
```

```
Neighbors: 3, established: 2
```

| | | |
|----------------------------|--------|-------------|
| | Sent | Received |
| Updates: | 3 | 15 |
| Notifications: | 0 | 0 |
| | Number | Memory Used |
| Attributes: | 12 | 1104 |
| AS Paths: | 10 | 400 |
| Communities: | 2 | 1080 |
| Extended communities: | 1 | 40 |
| Route Reflector Entries: | 0 | 0 |
| Route-map Cache Entries: | 0 | 0 |
| Filter-list Cache Entries: | 0 | 0 |

```

Next Hop Cache Entries:    2                80
Update messages queued:   0

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Main Table Version: 12
IGP notification: IGP notified

State: normal mode.
BGP Table Version: 12
Network Entries: 15, Soft Reconfig Entries: 0
Dampened Paths: 0, History Paths: 9

Prefixes:                  Allocated      Freed
Paths:                     15              0

Prefixes:                  Number         Memory Used
Paths:                     19              760

```

This table describes the significant fields shown in the display.

Table 26: show bgp process detail Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP is operating in | Indicates whether BGP is operating in standalone mode. |
| Autonomous System | Autonomous system number for the local system. |
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | Cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. |
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. |
| Enforce first AS enabled | Indicates that strict checking of the first autonomous system (AS) number in paths received from external BGP peers is enabled. |
| iBGP to IGP redistribution | Indicates internal redistribution is enabled using the bgp redistribution-internal command. |
| Treating missing MED as worst | Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled by the bgp bestpath med missing-as-worst command. |

| Field | Description |
|--|--|
| Always compare MED is enabled | Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the bgp bestpath med always command. |
| AS Path ignore is enabled | Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the bgp bestpath as-path ignore command. |
| Comparing MED from confederation peers | Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the bgp bestpath med confed command. |
| Comparing router ID for eBGP paths | Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the bgp bestpath compare-routerid command. |
| Default local preference | Default local preference value used for BGP routes. |
| Default keepalive | Default keepalive interval. This is controlled by the timers bgp command. |
| Graceful restart enabled | Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: <ul style="list-style-type: none"> • bgp graceful-restart • bgp graceful-restart purge-time • bgp graceful-restart stalepath-time • bgp graceful-restart restart-time • bgp graceful-restart graceful-reset |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the bgp scan-time command. |
| BGP Speaker Process | Speaker process responsible for receiving, processing and sending BGP messages. |
| Node | Node on which the specified process is executing. |
| Neighbors | Number of neighbors for which the specified process is responsible. |
| established | Number of neighbors that have connections in the established state for the specified process. |
| Updates | Number of update messages sent and received by the specified process. |
| Notifications | Number of notification messages sent and received by the specified process. |
| Attributes | Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information. |

| Field | Description |
|-------------------------|--|
| AS Paths | Number of unique autonomous system paths stored in the specified process and the amount of memory used by the AS path information. |
| Communities | Number of unique sets of community information stored in the specified process and the amount of memory used by them. |
| Extended communities | Number of unique sets of extended community information stored in the specified process and the amount of memory used by them. |
| Route Reflector Entries | Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them. |
| Nexthop Entries | Number of entries and memory usage for cached next- hop information. |
| Update messages queued | Total number of update messages queued to be sent across all neighbors for which the specified process is responsible. |
| Address family | Specified address family. |
| Dampening | Indicates whether dampening is enabled for the specified address family. |
| Client reflection | Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the bgp client-to-client reflection disable command. |
| Scan interval | Interval (in seconds) between BGP scans for the given address family. This is controlled by the bgp scan-time command. |
| Main Table Version | Last version of the local BGP database for the specified address family that was injected into the main routing table. |
| IGP notification | Indicates whether IGP has been notified of BGP convergence for the specified address family. |
| RIB has converged | Indicates whether the main routing table version has converged and the version at which it converged. |

| Field | Description |
|-------------------------------|---|
| State | <p>BGP system state for the specified address family and process. This may be one of the following:</p> <p>read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers.</p> <p>best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode.</p> <p>RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB.</p> <p>label allocation mode: Labels are allocated for the received prefixes based on the requirement.</p> <p>normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.</p> |
| BGP Table Version | Last version used in the BGP database for received routes. |
| Attribute download | Indicates whether the RIB attribute download is enabled. |
| Network Entries | Number of sets of prefix information held in the specified BGP process for the specified address family. |
| Soft Reconfig Entries | Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration. |
| Dampened Paths | Number of routes that are suppressed due to dampening for the specified address family. |
| History Paths | Number of routes that are currently withdrawn, but are being maintained to preserve dampening information. |
| Prefixes (Allocated/Freed) | Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Paths (Allocated/Freed) | Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Prefixes (Number/Memory Used) | Number of sets of prefix information currently allocated for the specified address family, and the amount of memory used by them. |
| Paths (Number/Memory Used) | Number of sets of route information currently allocated for the specified address family, and the amount of memory used by them. |

The following is sample output from the **show bgp process** command with the **performance-statistics** keyword:

```
RP/0/RP0/CPU0:router# show bgp process performance-statistics detail

BGP Speaker process: 0, Node: node0_0_CPU0
```

Restart count: 2
 Neighbors: 3, established: 2

| | Sent | Received |
|----------------|------|----------|
| Updates: | 20 | 20 |
| Notifications: | 0 | 0 |

| | Number | Memory Used |
|----------------------------|--------|-------------|
| Attributes: | 2 | 184 |
| AS Paths: | 2 | 48 |
| Communities: | 0 | 0 |
| Extended communities: | 0 | 0 |
| Route Reflector Entries: | 0 | 0 |
| Route-map Cache Entries: | 0 | 0 |
| Filter-list Cache Entries: | 0 | 0 |
| Next Hop Cache Entries: | 2 | 80 |
| Update messages queued: | 0 | |

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
 Read throttled 0 times
 Processed 14 inbound messages (time spent: 0.132 secs)
 Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
 Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
 Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
 Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
 Configuration: 1 requests (time spent: 0.002 secs)
 Operational data: 9 requests (time spent: 0.026 secs)

State: normal mode.
 BGP Table Version: 150
 Network Entries: 149, Soft Reconfig Entries: 0

| | Allocated | Freed |
|-----------|-----------|-------|
| Prefixes: | 149 | 0 |
| Paths: | 200 | 0 |

| | Number | Memory Used |
|-----------|--------|-------------|
| Prefixes: | 149 | 12516 |
| Paths: | 200 | 8000 |

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
 Scanner: 2 scanner runs (time spent: 0.008 secs)
 RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
 Process has converged for IPv4 Unicast.

First neighbor established: 1082604050s
 Entered DO_BESTPATH mode: 1082604055s
 Entered DO_RIBUPD mode: 1082604055s
 Entered Normal mode: 1082604055s
 Latest UPDATE sent: 1082604056s

This table describes the significant fields shown in the display.

Table 27: show bgp process performance-statistics Field Descriptions

| Field | Description |
|--|--|
| BGP is operating in | Indicates whether BGP is operating in standalone mode. |
| Autonomous system | Autonomous system number for the local system. |
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | The cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. |
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the bgp log neighbor changes disable command. |
| Enforce first AS enabled | Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled. |
| iBGP to IGP redistribution | Indicates internal redistribution is enabled using the bgp redistribution-internal command. |
| Treating missing MED as worst | Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled using the bgp bestpath med missing-as-worst command. |
| Always compare MED is enabled | Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This setting is controlled by the bgp bestpath med always command. |
| AS Path ignore is enabled | Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the bgp bestpath as-path ignore command. |
| Comparing MED from confederation peers | Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the bgp bestpath med confed command. |
| Comparing router ID for eBGP paths | Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the bgp bestpath compare-routerid command. |

| Field | Description |
|--------------------------|--|
| Default local preference | Default local preference value used for BGP routes. |
| Default keepalive | Default keepalive interval. This setting is controlled by the timers bgp command. |
| Graceful restart enabled | Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: bgp graceful-restart , bgp graceful-restart purge-time , bgp graceful-restart stalepath-time , bgp graceful-restart restart-time , and bgp graceful-restart graceful-reset . |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This setting is controlled by the bgp scan-time command in router configuration mode. |
| Address family | Specified address family. |
| Dampening | Indicates whether dampening is enabled for the specified address family. |
| Client reflection | Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the bgp client-to-client reflection disable command. |
| Scan interval | Interval (in seconds) between BGP scans for the given address family. This is controlled by the bgp scan-time command. |
| Main Table Version | Last version of the local BGP database for the specified address family that was injected into the main routing table. |
| IGP notification | Indicates whether IGP has been notified of BGP convergence for the specified address family. |
| Node | Node on which the process is executing. |
| Process | BGP process. |
| Speaker | Speaker process. The speaker process is responsible for receiving, processing and sending BGP messages. |
| Read | Real time (in seconds) spent reading messages from peers by this process. |
| Write | Real time (in seconds) spent writing messages to peers by this process. |
| Inbound | The real time (in seconds) spent processing messages read from peers by this process. |
| Config | Real time (in seconds) spent processing configuration commands by this process. |
| Data | Real time (in seconds) spent providing operational data by this process. |
| Conv | Indicates whether the process has converged after the initial update. |

| Field | Description |
|---|--|
| Nbr Estab | Time stamp (in seconds) recording the time when the first neighbor became established. |
| Bestpath | Time stamp (in seconds) recording the time the best-path calculation mode was entered. |
| RIB Inst | Time stamp (in seconds) recording the time RIB update mode was entered. |
| Read/Write | Time stamp (in seconds) recording the time normal mode was entered. |
| Last Upd | Time stamp (in seconds) recording the time the last update was sent to a neighbor. |
| Address Family IPv4 Unicast converged in <i>n</i> seconds | Indicates that BGP has reached initial convergence for the IPv4 unicast address family. The time taken for convergence is shown. |

The following is sample output from the **show bgp process** command with the **performance-statistics** and **detail** keywords:

```
RP/0/RP0/CPU0:router# show bgp process performance-statistics detail
```

```
BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2
```

| | Sent | Received |
|----------------|------|----------|
| Updates: | 20 | 20 |
| Notifications: | 0 | 0 |

| | Number | Memory Used |
|----------------------------|--------|-------------|
| Attributes: | 2 | 184 |
| AS Paths: | 2 | 48 |
| Communities: | 0 | 0 |
| Extended communities: | 0 | 0 |
| Route Reflector Entries: | 0 | 0 |
| Route-map Cache Entries: | 0 | 0 |
| Filter-list Cache Entries: | 0 | 0 |
| Next Hop Cache Entries: | 2 | 80 |
| Update messages queued: | 0 | |

```
Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)
```

```
State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0
```

| Allocated | Freed |
|-----------|-------|
|-----------|-------|

```

Prefixes:                149                0
Paths:                   200                0

                          Number            Memory Used
Prefixes:                149            12516
Paths:                   200            8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s

```

This table describes the significant fields shown in the display.

Table 28: show bgp process performance-statistics detail Field Descriptions

| Field | Description |
|---------------------------|--|
| Process | The specified process. |
| Location | Node in which the specified process is executing. |
| Neighbors | Number of neighbors for which the specified process is responsible. |
| established | Number of neighbors that have connections in the established state for the specified process. |
| Updates | Number of update messages sent and received by the specified process. |
| Notifications | Number of notification messages sent and received by the specified process. |
| Attributes | Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information. |
| AS Paths | Number of unique autonomous system paths stored in the specified process, and the amount of memory used by the AS path information. |
| Communities | Number of unique sets of community information stored in the specified process and the amount of memory used by them. |
| Extended communities | Number of unique sets of extended community information stored in the specified process and the amount of memory used by them. |
| Route Reflector Entries | Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them. |
| Route-map Cache Entries | Number of entries and memory usage for cached results for applying a route map. |
| Filter-list Cache Entries | Number of entries and memory usage for cached results for applying an AS path filter list. |

| Field | Description |
|------------------------|---|
| Next Hop Cache Entries | Number of entries and memory usage for cached next-hop information. |
| Update messages queued | Number of update messages queued to be sent across all neighbors for which the specified process is responsible. |
| Read | Indicates the number of messages read by the process, the total size of read messages, the number of read operations performed, and the real time spent by the process performing read operations. |
| Read throttled | Number of times that reading from TCP has been throttled due to a backlog of messages read but not processed. |
| inbound messages | Number of read messages that have been processed and the real time spent processing inbound messages. |
| Wrote | Amount of data that has been written by the process, the number of write operations performed, and the real time spent by the process performing write operations. |
| Processing write list | Number of messages written from write lists, the number of times the write list has been processed, and the real time spent processing the write list. Note Write lists typically contain only update messages. |
| Processing write queue | Number of messages written from write queues, number of times the write queue has been processed, and the real time spent processing the write queue. |
| Socket setup | Number of socket setup operations performed and the real time spent during socket setup operations. |
| Configuration | Number of configuration requests received by the process and the real time spent processing configuration requests. |
| Operational data | Number of requests for operational data (for show commands) received by the process and the real time spent processing operation data requests |

| Field | Description |
|-------------------------------|---|
| State | <p>BGP system state for the specified address family and process. This may be one of the following:</p> <p>read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers.</p> <p>best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode.</p> <p>RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB.</p> <p>label allocation mode: Labels are allocated for the received prefixes based on the requirement.</p> <p>normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.</p> |
| BGP Table Version | Last version used in the BGP database for received routes. |
| Network Entries | Number of sets of prefix information held in the specified BGP process for the specified address family. |
| Soft Reconfig Entries | Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration. |
| Dampened Paths | Number of routes that are suppressed due to dampening for the specified address family. |
| History Paths | Number of routes that are currently withdrawn, but are being maintained to preserve dampening information. |
| Prefixes (Allocated/Freed) | Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Paths (Allocated/Freed) | Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Prefixes (Number/Memory Used) | Number of sets of prefix information currently allocated for the specified address family and amount of memory used by them. |
| Paths (Number/Memory Used) | Number of sets of route information currently allocated for the specified address family and amount of memory used by them. |
| Updates generated | Number of prefixes for which updates have been generated, the number of messages used to advertise the updates, the number of update generation runs performed, and the real time spent generating updates for the specified address family. |
| Scanner | Number of times the scanner has run for the specified address family and real time spent in scanner processing. |

| Field | Description |
|----------------------------|--|
| RIB Update | Number of global routing information base update runs performed for the specified address family, number of prefixes installed, withdrawn, or modified in the global RIB during these runs, and real time spent performing these runs. |
| Process has converged | Indicates whether the process has reached initial convergence for the specified address family. |
| First neighbor established | Time stamp (in seconds) recording the time the first neighbor in the process was established. |
| Entered DO_BESTPATH mode | Time stamp (in seconds) recording the time best-path calculation mode was entered. |
| Entered DO_RIBUPD mode | Time stamp (in seconds) recording the time RIB update mode was entered. |
| Entered Normal mode | Time stamp (in seconds) recording the time normal mode was entered. |
| Last UPDATE sent | Time stamp (in seconds) recording the time the last update was sent to a neighbor. |

The following is sample output from show bgp process detail command with information on additional paths send and receive information:

```

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASDOT
Autonomous System: 100
Router ID: 22.22.22.22 (manually configured)
Default Cluster ID: 2.2.2.2 (manually configured)
Active Cluster IDs: 2.2.2.2
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 120
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

.....
.....

Prefixes:           Allocated      Freed
Paths:              12              0
Path-elems:         60              0
Path-elems:         12              0

                          Number      Mem Used
Prefixes:              12          1200
Paths:                  60          3120
Path-elems:             12          624

```

Related Commands

| Command | Description |
|---|---|
| bgp bestpath as-path ignore, on page 30 | Sets the autonomous system path length to ignore when calculating preferred paths. |
| bgp bestpath compare-routerid, on page 32 | Compare identical routes received from external BGP (eBGP) peers during the best-path selection process and select the route with the lowest router ID. |
| bgp bestpath med always, on page 34 | Compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. |
| bgp bestpath med missing-as-worst, on page 38 | Assume paths with no MED attribute have the most undesirable MED value possible when performing path selection. |
| bgp cluster-id, on page 46 | Enables reflection of routes between route reflector clients using a BGP route reflector. |
| bgp cluster-id, on page 46 | Configure the cluster ID if the BGP cluster has more than one route reflector. |
| bgp default local-preference, on page 54 | Sets the default local preference value. |
| bgp redistribute-internal, on page 70 | Allows the redistribution of iBGP routes into an IGP such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF). |
| bgp router-id, on page 72 | Configures a fixed router ID for a BGP-speaking router. |
| default-metric (BGP), on page 108 | Sets default metric values for the BGP. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| bgp scan-time, on page 73 | Configures scanning intervals. |
| timers bgp, on page 353 | Sets default BGP timers. |

show bgp regexp

To display routes matching the autonomous system path regular expression, use the **show bgp regexp** command in XR EXEC mode.

show bgp regexp *regular-expression*

| | | |
|---------------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | <i>regular-expression</i> | Regular expression to match the BGP autonomous system paths. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp regexp** command to display all routes in the specified BGP table whose autonomous system path is matched by the specified regular expression.



Note If the regular expression contains spaces and parentheses, it must be specified and surrounded by quotation marks.

Task ID

Task ID Operations

bgp read

Examples

The following is sample output from the **show bgp regexp** command:

```
RP/0/RP0/CPU0:router# show bgp regexp "^3 "

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 64
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*>i172.20.17.121    10.0.101.2                100      0 3 2000 3000 i
*>i10.0.0.0         10.0.101.2                100      0 3 100 1000 i
*>i172.5.23.0/24    10.0.101.2                100      0 3 4 60 4378 i
```

This table describes the significant fields shown in the display.

Table 29: show bgp regexp Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is placed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP address of a network entity. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|--|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |
| show bgp route-policy, on page 312 | Displays BGP information about networks that match an outbound route policy. |

show bgp route-policy

To display Border Gateway Protocol (BGP) information about networks that match an outbound route policy, use the **show bgp route-policy** command in XR EXEC mode.

show bgp route-policy *route-policy-name* [**standby**]

| | | |
|--------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | <i>route-policy-name</i> | Name of a route policy. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

A route policy must be configured to use this command. When the **show bgp route-policy** command is entered, routes in the specified BGP table are compared with the specified route policy, and all routes passed by the route policy are displayed.

If a pass clause is encountered while the route policy is being applied to the route and the route policy processing completes without hitting a drop clause, the route is displayed. The route is not displayed if a drop clause is encountered, if the route policy processing completes without hitting a pass clause, or if the specified route policy does not exist.

The information displayed does not reflect modifications the policy might make to the route. To display such modifications, use the **show bgp policy** command.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp route-policy** command in XR EXEC mode:

```
RP/0/RP0/CPU0:router# show bgp route-policy p1

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 729
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*  10.13.0.0/16    192.168.40.24                0 1878 704 701 200 ?
*  10.16.0.0/16    192.168.40.24                0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 30: show bgp route-policy Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|---|
| aggregate-address, on page 12 | Configures an aggregate entry in a BGP routing table. |
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor |
| route-policy | Configures a route policy. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp policy, on page 283 | Displays advertisements under a proposed policy. |

show bgp session-group

To display information about the Border Gateway Protocol (BGP) configuration for session groups, use the **show bgp session-group** command in XR EXEC mode.

show bgp session-group *group-name* {**configuration** [**defaults**] [**nvgen**] | **inheritance** | **users**}

Syntax Description

| | |
|----------------------|---|
| <i>group-name</i> | Name of the session family group to display. |
| configuration | (Optional) Displays the effective configuration for the session group, including any inherited configuration. |
| defaults | (Optional) Displays all configuration, including default configuration. |
| nvgen | (Optional) Displays output in the form of the show running-config command. If the defaults keyword also is specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | (Optional) Displays the session groups from which this session group inherits configuration. |
| users | (Optional) Display the session groups, neighbor groups, and neighbors that inherit configuration from this session group. |

Command Default

No default behavior or value

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp session-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a session group, including any configuration inherited from other session groups through application of the **use** command. The source for each configured command is also displayed.

Use the **defaults** keyword to display the value of all configuration, including default configuration. Use the **nvgen** keyword to display configuration in the form of the **show running-config** command output. Output in this form is suitable for cutting and pasting into a configuration session.

Use the **show bgp session-group** command with the *group-name* **inheritance** argument and keyword to display the session groups from which the specified session group inherits configuration.

Use the **show bgp session-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and session groups that inherit configuration from the specified session group.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

For the example shown here, the following configuration is used:

```
session-group group3
  advertisement-interval 5
  dmzlink-bw
!
session-group group1
  use session-group group2
  update-source Loopback0
!
session-group group2
  use session-group group3
  ebgp-multihop 2
```

The following example shows the **show bgp session-group** command with the **configuration** keyword:

```
RP/0/RP0/CPU0:router# show bgp session-group group1 configuration

session-group group1
  advertisement-interval 5[s:group2 s:group3]
  ebgp-multihop 2 [s:group2]
  update-source Loopback0 []
  dmzlink-bandwidth [s:group2 s:group3]
```

The source of each command is shown to the right of the command. For example, **update-source** is configured directly on session group group1. The **dmzlink-bandwidth** command is inherited from session group group2, which in turn inherits it from session group group3.

The following example shows the **show bgp session-group** command with the **users** keyword:

```
RP/0/RP0/CPU0:router# show bgp session-group group2 users

IPv4 Unicast:a:group1
```

The following example shows the **show bgp session-group** command with the **inheritance** keyword.

```
RP/0/RP0/CPU0:router# show bgp session-group group1 inheritance

Session:s:group2 s:group3
```

The command output shows that the session group group1 directly uses the group2 session group. The group2 session group uses the group3 session group.

This table describes the significant fields shown in the display.

Table 31: show bgp session-group Field Descriptions

| Field | Description |
|-----------|--|
| [] | Configures the command directly on the specified session group. |
| s: | Indicates the name that follows is a session group. |
| a: | Indicates the name that follows is an address family group. |
| n: | Indicates the name that follows is a neighbor group. |
| [dflt] | Indicates the command is not explicitly configured or inherited, and the default value for the command is used. This field may be shown when the defaults keyword is specified. |
| <not set> | Indicates that the default is for the command to be disabled. This field may be shown when the defaults keyword is specified. |

Related Commands

| Command | Description |
|--|---|
| session-group, on page 189 | Configures a BGP session group. |
| show bgp neighbor-group, on page 252 | Displays information about the BGP configuration for neighbor groups. |
| show bgp neighbors, on page 256 | Displays information about BGP connections to neighbors. |

show bgp sessions

To display brief information about BGP neighbors, use the **show bgp sessions** command in XR EXEC mode.

show bgp sessions [**not-established**] [**not-nsr-ready**]

| | | |
|--------------------|---|--|
| Syntax Description | not-established (Optional) Displays all the neighbors that are not in established state | |
| | not-nsr-ready (Optional) Displays all the neighbors that are not nonstop routing (NSR) ready. | |
| Command Default | No default behavior or values | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | The show bgp sessions command with the not-established keyword shows BGP peers which are yet to establish their peering relationship. | |
| | The show bgp session command with the and not-nsr-ready keyword shows BGP peers which are yet to reach the nsr ready state. | |
| Task ID | Task ID | Operations |
| | bgp | read |
| Related Commands | Command | Description |
| | show bgp neighbors, on page 256 | Displays information about Border Gateway Protocol (BGP) connections to neighbors. |

show bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the **show bgp summary** command in XR EXEC mode.

```
show bgp [{ipv4 {unicast | labeled-unicast | all} | ipv6 {unicast | all | labeled-unicast}}] summary [standby]
```

| | | |
|---------------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Use the **show bgp summary** command to display a summary of the neighbors for which the specified address family and subaddress family are enabled. If the neighbor does not have the specified address family and subaddress family enabled, it is not included in the output of the **show** command. If the **all** keyword is specified for the address family or subaddress family, a summary for each combination of address family and subaddress family is displayed in turn.

The table versions shown in the output (RecvTblVer, bRIB/RIB, SendTblVer, and TblVer) are specific to the specified address family and subaddress family. All other information is global.

The table versions provide an indication of whether BGP is up to date with all work for the specified address family and subaddress family.

- bRIB/RIB < RecvTblVer—Some received routes have not yet been considered for installation in the global routing table.
- TblVer < SendTblVer—Some received routes have been installed in the global routing table but have not yet been considered for advertisement to this neighbor.

Task ID

Task ID Operations

bgp read

Examples

The following is sample output from the **show bgp summary** command:

```
RP/0/RP0/CPU0:router#show bgp summary
```

```
BGP router identifier 10.0.0.0, local AS number 2
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 1
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

| Process | RecvTblVer | bRIB/RIB | LabelVer | ImportVer | SendTblVer |
|---------|------------|----------|----------|-----------|------------|
| Speaker | 1 | 0 | 1 | 1 | 0 |

| Neighbor | Spk | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down | St/PfxRcd |
|------------|-----|----|---------|---------|--------|-----|------|----------|-----------|
| 10.0.101.0 | 0 | 2 | 0 | 0 | 0 | 0 | 0 | 00:00:00 | Idle |
| 10.0.101.1 | 0 | 2 | 0 | 0 | 0 | 0 | 0 | 00:00:00 | Idle |

This table describes the significant fields shown in the display.

Table 32: show bgp summary Field Descriptions

| Field | Description |
|---------------------------|--|
| BGP router identifier | IP address of the router. |
| local AS number | Autonomous system number set by the router bgp, on page 182 command. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP table state | State of the BGP database. |

| Field | Description |
|---|--|
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was injected into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| BGP is operating in | Specifies BGP is operating in standalone mode. |
| Process | BGP process. |
| RecvTblVer | Last version used in the BGP database for received routes. |
| bRIB/RIB | Last version of the local BGP database that was injected into the main routing table. |
| LabelVer | Label version used in the BGP database for label allocation. |
| ImportVer | Last version of the local BGP database for importing routes. |
| SendTblVer | Latest version of the local BGP database that is ready to be advertised to neighbors. |
| Some configured eBGP neighbors do not have any policy | Some external neighbors exist that do not have both an inbound and outbound policy configured for every address family, using the route-policy (BGP) command. In this case, no prefixes are accepted and advertised to those neighbors. |
| Neighbor | IP address of a neighbor. |
| Spr | Speaker process that is responsible for the neighbor. Always 0. |
| AS | Autonomous system. |
| MsgRcvd | Number of BGP messages received from a neighbor. |
| MsgSent | Number of BGP messages sent to a neighbor. |
| TblVer | Last version of the BGP database that was sent to a neighbor. |
| InQ | Number of messages from a neighbor waiting to be processed. |
| OutQ | Number of messages waiting to be sent to a neighbor. |
| Up/Down | Length of time in (hh:mm:ss) that the BGP session has been in Established state, or the time since the session left Established state, if it is not established. |

| Field | Description |
|-----------|--|
| St/PfxRcd | <p>If the BGP session is not established, the current state of the session. If the session is established, the number of prefixes the router has received from the neighbor.</p> <p>If the number of prefixes received exceeds the maximum allowed (as set by the maximum-prefix command), “(PfxRcd)” appears.</p> <p>If the connection has been shut down using the shutdown command, “(Admin)” appears.</p> <p>If the neighbor is external and it does not have an inbound and outbound policy configured for every address family, an exclamation mark (!) is inserted at the end of the state when using the route-policy (BGP) command.</p> <p>If the connection has been shut down due to out of memory (OOM), “(OOM)” appears.</p> |

Related Commands

| Command | Description |
|---|--|
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |

show bgp summary nsr

To display the summary of Border Gateway Protocol (BGP) neighbor state and nonstop routing (NSR) state information, use the **show bgp summary nsr** command in XR EXEC mode.

show bgp summary nsr [**standby**]

| | | |
|---------------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.2.1 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operations |
| | bgp | read |

Examples

The following is sample output from the **show bgp summary nsr** command:

```
RP/0/RP0/CPU0:router# show bgp summary nsr

BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 11034 (Reached)
```

BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

node0_1_CPU0 Speaker

```
Entered mode Standby Ready           : Feb  3 14:22:00
Entered mode TCP NSR Setup           : Feb  3 14:22:00
Entered mode TCP NSR Setup Done      : Feb  3 14:22:01
Entered mode TCP Initial Sync        : Feb  3 14:22:01
Entered mode TCP Initial Sync Done   : Feb  3 14:22:44
Entered mode FPBSN processing done   : Feb  3 14:22:44
Entered mode Update processing done  : Feb  3 14:22:44
Entered mode BGP Initial Sync        : Feb  3 14:22:44
Entered mode BGP Initial Sync done   : Feb  3 14:22:49
Entered mode NSR Ready               : Feb  3 14:22:49
```

Current BGP NSR state - NSR Ready achieved at: Feb 3 14:22:49

NSR State READY notified to Redcon at: Feb 4 07:44:43

| Process | RcvTblVer | bRIB/RIB | LabelVer | ImportVer | SendTblVer | StandbyVer |
|---------|-----------|----------|----------|-----------|------------|------------|
| Speaker | 13037 | 13037 | 13037 | 13037 | 13037 | 13037 |

| Neighbor | Spk | AS | TblVer | SyncVer | AckVer | NBRState | NSRState |
|------------|-----|-----|--------|---------|--------|-------------|-----------|
| 2.2.2.2 | 0 | 302 | 13037 | 13037 | 13037 | Established | NSR Ready |
| 10.0.101.5 | 0 | 100 | 13037 | 13037 | 13037 | Established | NSR Ready |

The following example shows sample output from the **show bgp summary nsr** command with the **standby** keyword:

RP/0/RP0/CPU0:router# **show bgp summary nsr standby**

```
BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 0 (Not Reached)
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

node0_0_CPU0 Speaker

```
Entered mode Standby Ready           : Feb  3 14:22:03
Entered mode TCP Replication          : Feb  3 14:22:03
Entered mode TCP Init Sync Done      : Feb  3 14:22:47
Entered mode NSR Ready               : Feb  3 14:22:52
```

| Process | RcvTblVer | bRIB/RIB | LabelVer | ImportVer | SendTblVer | StandbyVer |
|---------|-----------|----------|----------|-----------|------------|------------|
| Speaker | 13037 | 0 | 0 | 13037 | 0 | 0 |

| Neighbor | Spk | AS | TblVer | SyncVer | AckVer | NBRState | NSRState |
|------------|-----|-----|--------|---------|--------|-------------|-----------|
| 2.2.2.2 | 0 | 302 | 13037 | 0 | 1 | Established | NSR Ready |
| 10.0.101.5 | 0 | 100 | 13037 | 0 | 1 | Established | NSR Ready |

This table describes the significant fields shown in the display.

Table 33: show bgp summary nsr Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | IP address of the router. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| Non-stop routing | State of the Nonstop routing. |
| BGP table state | State of the BGP database. |
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was injected into the main routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| BGP is operating in | Specifies BGP is operating in standalone mode. |
| Entered mode | The successive transition of various states of TCP and BGP, leading to the NSR ready state. Note This is used for monitoring and debugging purposes. |
| SyncVer | The version which has synced to standby for this neighbor. |
| AckVer | The version which the neighbor has acknowledge. |
| NBRState | State of the BGP neighbor. |
| NSRState | Neighbor NSR state. |

Related Commands

| Command | Description |
|---|---|
| nsr (BGP), on page 155 | Activates Border Gateway Protocol (BGP) nonstop routing (NSR) |
| show bgp nsr, on page 277 | Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information. |

show bgp table

To display the status of all Border Gateway Protocol (BGP) neighbors for a particular Address Family (AF) in the global address table, use the **show bgp table** command in XR EXEC mode.

show bgp table

| | | |
|--------------------|--|---|
| Syntax Description | ipv4 (Optional) Specifies IP Version 4 address prefixes. | |
| | unicast (Optional) Specifies unicast address prefixes. | |
| | ipv6 (Optional) Specifies IP Version 6 address prefixes. | |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | The set default-afi command is used to specify the default address family for the session, and the set default-safi command is used to specify the default subaddress family for the session. See <i>Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference</i> for detailed information and syntax for the set default-afi and set default-safi commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast. | |
| | Use the show bgp table command to display a brief summary of the neighbors for which the specified address family (AFI) and subaddress family (SAFI) are enabled. If the AFI and/or SAFI is not enabled, the command will only display the column headings. | |
| Task ID | Task ID | Operations |
| | bgp | read |
| Related Commands | Command | Description |
| | show bgp neighbor-group, on page 252 | Displays information about the Border Gateway Protocol (BGP) configuration for neighbor groups. |
| | show bgp neighbors, on page 256 | Displays information about Border Gateway Protocol (BGP) connections to neighbors. |

| Command | Description |
|---|---|
| show bgp summary, on page 320 | Displays the status of all Border Gateway Protocol (BGP) connections. |

show bgp truncated-communities

To display routes in the Border Gateway Protocol (BGP) routing table for which inbound policy or aggregation has exceeded the maximum number of communities that may be attached, use the **show bgp truncated-communities** command in XR EXEC mode.

show bgptruncated-communities standby

| | | |
|--------------------|---|---|
| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| | unicast | (Optional) Specifies unicast address prefixes. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| | ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| | all | (Optional) For address family, specifies prefixes for all address families. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

Use the **show bgp truncated-communities** command to display those routes in the specified BGP routing table in which the buffers used to store communities or extended communities have overflowed. An overflow occurs if an attempt is made to associate more communities or extended communities with the route than fits in a BGP update message. This can happen due to modification of communities or extended communities during aggregation or when inbound policy is applied.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples

The following is sample output from the **show bgp truncated-communities** command:

```
RP/0/RP0/CPU0:router# show bgp truncated-communities

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 3042
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*  10.13.0.0/16    192.168.40.24                      0 1878 704 701 200 ?
*> 10.16.0.0/16    192.168.40.24                      0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 34: show bgp truncated-communities Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP Identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|---|
| aggregate-address, on page 12 | Creates an aggregate entry in a BGP routing table. |
| network (BGP), on page 143 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp, on page 193 | Displays entries in the BGP routing table. |

show bgp update-group

To display Border Gateway Protocol (BGP) information for update groups, use the **show bgp update-group** command in XR EXEC mode.

```
show bgp [{ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} [{ipv4 {unicast | labeled-unicast} | ipv6 unicast}] | vpnv6 unicast}] update-group [{neighbor ip-address | process-id . index [{summary | performance-statistics}]]]
show bgp [{ipv4 {unicast | labeled-unicast | all} | ipv6 {unicast | all | labeled-unicast}}] update-group [{neighbor ip-address | process-id . index [{summary | performance-statistics}]]] [standby]
```

| Syntax Description | ipv4 | (Optional) Specifies IP Version 4 update groups. |
|--------------------|--|---|
| | unicast | (Optional) Specifies unicast update groups. |
| | labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| | all | (Optional) Displays both unicast update groups. |
| | ipv6 | (Optional) Specifies IP Version 6 update groups. |
| | all | (Optional) Displays both IP Version 4 and IP Version 6 update groups. |
| | neighbor <i>ip-address</i> | (Optional) Specifies information on an update group for a specific neighbor. |
| | <i>process-id.index</i> | (Optional) Update group index. Process ID range is 0 to 254. Index range is 0 to 4294967295. The <i>process id.index</i> argument is specified as follows: process ID (dot) index. In standalone mode, the process ID is always 0. |
| | summary | (Optional) Specifies summary of update group members. |
| | performance-statistics | (Optional) Specifies performance information about the updates generated for the update group. |
| | standby | (Optional) Displays information about the standby card. |
| Command Default | If no address family or subaddress family is specified, the default address family and subaddress family specified using the set default-afi and set default-safi commands are used. | |
| Command Modes | XR EXECEXEC | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| | Release 7.1.1 | The command output was modified to include the performance statistics for BGP peer that was not advertised.. |

Usage Guidelines



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *System Management Command Reference for Cisco NCS 6000 Series Routers* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Every BGP neighbor is automatically assigned to an update group for each address family that is enabled on the neighbor. Neighbors that have similar outbound policy, such that they are sent the same updates, are placed in the same update group.

Use the **show bgp update-group** command to display the update groups and a list of the neighbors that belong to the update group.

Use the **show bgp update-group neighbor** command to display details about the update group to which a neighbor belongs for the specified address family.

Use the **summary** keyword to display a summary of the neighbors belonging to the specified update group. The display format is the same as for the [show bgp summary, on page 320](#) command.

Use the **performance-statistics** keyword to display information about the number of prefixes processed and the time taken to generate updates for the specified update group.



Note Update group indexes are not necessarily persistent over a process restart. If a BGP process restarts, the index of the update group to which a particular neighbor is assigned may be different, though the set of neighbors belonging to the update group is the same.

Task ID

| Task ID | Operations |
|---------|------------|
|---------|------------|

| | |
|-----|------|
| bgp | read |
|-----|------|

This is sample output from the **show bgp update-group performance-statistics** command that shows the total numbers of prefixes not advertising in any AFI or SAFI, including repeating counts on 1 or more prefixes:

```
Router# show bgp update-group performance-statistics
```

```
Update group for IPv4 Unicast, index 0.1:
```

```
..
```

```
Update timer last processed: Sep 23 00:10:15.350
```

```
Not-Advertised Stats:
```

```
Non-Client to Non-Client      : 105      Sep 23 00:10:15.350
Path Not Available             : 132      Sep 23 00:10:15.350
```

Related Commands

| Command | Description |
|---|---|
| maximum-prefix (BGP), on page 134 | Limits the number of prefixes that can be received from a neighbor. |

| Command | Description |
|---|--|
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp summary, on page 320 | Displays the status of all BGP connections. |
| shutdown (BGP), on page 339 | Disables a neighbor without removing its configuration. |

show protocols (BGP)

To display information about the Border Gateway Protocol (BGP) instances running on the router, use the **show protocols** command in XR EXEC mode and specify either the **bgp** or **all** keyword.

show protocols [{**ipv4** | **ipv6** | **afi-all**}] [{**allprotocol**}]

Syntax Description

| | |
|-----------------|---|
| ipv4 | (Optional) Specifies the IP Version 4 address family. |
| ipv6 | (Optional) Specifies the IP Version 6 address family. |
| afi-all | (Optional) Specifies all address families. |
| all | (Optional) Specifies all protocols for a given address family. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are bgp , isis , rip , eigrp , and ospf . For the IPv6 address family, the options are bgp , eigrp , isis , and ospfv3 . |

Command Default

Default is IPv4.

Command Modes

XR EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show protocols** command to get information about the protocols running on the router and to quickly determine which protocols are active. The command is designed to summarize the important characteristics of the running protocol, and command output varies depending on the specific protocol selected. For BGP, the command output lists the protocol ID, peers with elapsed time since last reset, and miscellaneous information, such as external and internal local distances and sourced routes.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |
| rib | read |

Examples

The following example shows the display for the **show protocols** command using the **bgp** keyword:


```

RP/0/RP0/CPU0:router# show protocols bgp

Routing Protocol "BGP 40"

Address Family IPv4 Unicast:
  Distance: external 20 internal 200 local 200
  Sourced Networks:
    10.100.0.0/16 backdoor
    10.100.1.0/24
    10.100.2.0/24
  Routing Information Sources:
    Neighbor      State/Last update received
    10.5.0.2      Idle
    10.9.0.3      Idle

```

This table describes the significant fields shown in the display.

Table 35: show protocols (BGP) Field Descriptions

| Field | Description |
|-----------------------------|--|
| Routing Protocol: | Identifies BGP as the running protocol and displays the BGP AS number. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| Address Family | Specifies the address family. This can be IPv4 Unicast or IPv6 Unicast. |
| Distance: external | Specifies the distance BGP sets when installing eBGP routes into the RIB. eBGP routes are routes received from eBGP peers. The RIB uses the distance as a tiebreaker when several protocols install a route for the same prefix. |
| Distance: internal | Specifies the distance BGP sets for routes received from iBGP peers. |
| Distance: local | Specifies the distance BGP sets for locally generated aggregates and backdoor routes. |
| Sourced Networks | List of locally sourced networks. These are networks sourced using the network command. |
| Routing information Sources | List of configured BGP neighbors. |
| Neighbor | Address of a BGP neighbor. |
| State/Last update received | State of each neighbor and the time since the last update was received from the neighbor if it is established. |

show tcp brief

To display a concise description of TCP connection endpoints, use the **show tcp brief** command in the XR EXEC mode.

show tcp brief [*location node-id*]

| Syntax Description | location <i>node-id</i> (Optional) Specifies location information for the specified node ID. The node ID variable is mentioned in the <i>rack/slot/module</i> notation. | | | | |
|---------------------------|--|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | XR EXEC mode | | | | |
| Command History | <table> <tr> <th>Release</th><th>Modification</th></tr> <tr> <td>Release 3.7.2</td><td>This command was introduced.</td></tr> </table> | Release | Modification | Release 3.7.2 | This command was introduced. |
| Release | Modification | | | | |
| Release 3.7.2 | This command was introduced. | | | | |
| Usage Guidelines | No specific guidelines impact the use of this command. | | | | |
| Task ID | <table> <tr> <th>Task ID</th><th>Operation</th></tr> <tr> <td>bgp</td><td>read</td></tr> </table> | Task ID | Operation | bgp | read |
| Task ID | Operation | | | | |
| bgp | read | | | | |

Example

The following is a sample output from the **show tcp brief** command:

```
RP/0/0/CPU0:ios#show tcp brief
```

| PCB | VRF-ID | Recv-Q | Send-Q | Local Address | Foreign Address | State |
|------------|------------|--------|--------|----------------|-----------------|--------|
| 0x08789b28 | 0x60000000 | 0 | 0 | :::179 | :::0 | LISTEN |
| 0x08786160 | 0x00000000 | 0 | 0 | :::179 | :::0 | LISTEN |
| 0xecb0c9f8 | 0x60000000 | 0 | 0 | 10.0.0.1:12404 | 10.0.0.2:179 | ESTAB |
| 0x0878b168 | 0x60000000 | 0 | 0 | 11.0.0.1:179 | 11.0.0.2:61177 | ESTAB |
| 0xecb0c6b8 | 0x60000000 | 0 | 0 | 0.0.0.0:179 | 0.0.0.0:0 | LISTEN |
| 0x08781590 | 0x00000000 | 0 | 0 | 0.0.0.0:179 | 0.0.0.0:0 | LISTEN |

shutdown (BGP)

To disable a neighbor without removing its configuration, use the **shutdown** command in an appropriate configuration mode. To re-enable the neighbor and reestablish a Border Gateway Protocol (BGP) session, use the **no** form of this command.

shutdown [**inheritance-disable**]
no shutdown [**inheritance-disable**]

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | inheritance-disable (Optional) Overrides the value of a shutdown command inherited from a neighbor group or session group. | |
| Command Default | Neighbors are not shutdown. | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the shutdown command to terminate any active session for the specified neighbor and remove all associated routing information. Use of the shutdown command with a neighbor group or session group may suddenly terminate a large number of BGP neighbor sessions because all neighbors using the neighbor group or session group may be affected.</p> <p>Use the show bgp summary command to display a summary of BGP neighbors. Neighbors that are idle due to the shutdown command are displayed with the “Idle (Admin)” state.</p> <p>If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.</p> | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | The following example shows that any active session for neighbor 192.168.40.24 is disabled: | |

```
RP/0/RP0/CPU0:router(config)# router bgp 1
```

```
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

In the following example, the session remains active for neighbor 192.168.40.24 because the inherited **shutdown** command has been overridden:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# shutdown
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RP0/CPU0:router(config-bgp-nbr)# shutdown inheritance-disable
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```


Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| show bgp summary, on page 320 | Displays the status of all BGP connections. |

socket receive-buffer-size

To set the size of the receive buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket receive-buffer-size** command in an appropriate configuration mode. To set the size of the receive buffers to the default size, use the **no** form of this command.

socket receive-buffer-size *socket-size* [*bgp-size*]
no socket receive-buffer-size [*socket-size*] [*bgp-size*]

| | | |
|--|---|------------------------------|
| Syntax Description | <i>socket-size</i> Size (in bytes) of the receive-side socket buffers. Range is 512 to 131072. | |
| | <i>bgp-size</i> (Optional) Size (in bytes) of the receive buffers in BGP. Range is 512 to 131072. | |
| Command Default | <i>socket-size</i> : 32,768 bytes <i>bgp-size</i> : 4,032 bytes | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the socket receive-buffer-size command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router. | |
|  Note | Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely. | |
| | Use the receive-buffer-size command on individual neighbors to change the values set by the socket receive-buffer-size command. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | The following example shows how to set the receive buffer sizes for all neighbors to 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer: | |

```
RP/0/RP0/CPU0:router(config)# router bgp 1  
RP/0/RP0/CPU0:router(config-bgp)# socket receive-buffer-size 65536 8192
```


Related Commands

| Command | Description |
|--|--|
| receive-buffer-size, on page 167 | Sets the size of the receive buffers for a BGP neighbor. |
| socket send-buffer-size, on page 343 | Sets the size of the send buffers for all BGP neighbors. |

socket send-buffer-size

To set the size of the send buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default size, use the **no** form of this command.

socket send-buffer-size *socket-size* [*bgp-size*]
no socket send-buffer-size [*socket-size*] [*bgp-size*]

| | | |
|--|---|------------------------------|
| Syntax Description | <i>socket-size</i> Size (in bytes) of the send-side socket buffers. Range is 4096 to 131072. | |
| | <i>bgp-size</i> (Optional) Size (in bytes) of the send buffers in BGP. Range is 4096 to 131072. | |
| Command Default | <i>socket-size</i> : 10240 bytes <i>bgp-size</i> : 4096 bytes | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the socket send-buffer-size command to increase the buffer size when sending updates to neighbors. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router. | |
|  Note | Increasing the socket buffer size uses more memory only when more messages are waiting to be sent by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely. | |
| | Use the send-buffer-size command on individual neighbors to change the values set by the socket send-buffer-size command. | |
| Task ID | Task ID | Operations |
| | bgp | read, write |
| Examples | The following example shows how to set the send buffer sizes for all neighbors to 8192 bytes for the socket buffer and the BGP buffer: | |

```
RP/0/RP0/CPU0:router(config)# router bgp 1  
RP/0/RP0/CPU0:router(config-bgp)# socket send-buffer-size 8192 8192
```


Related Commands

| Command | Description |
|---|---|
| send-buffer-size, on page 183 | Sets the size of the send buffers for a BGP neighbor. |
| socket receive-buffer-size, on page 341 | Sets the size of the receive buffers for all BGP neighbors. |

soft-reconfiguration inbound

To configure the software to store updates received from a neighbor, use the **soft-reconfiguration inbound** command in an appropriate configuration mode. To disable storing received updates, use the **no** form of this command.

soft-reconfiguration inbound [{always | inheritance-disable}]
no soft-reconfiguration inbound [{always | inheritance-disable}]

| | | |
|---|---|------------------------------|
| Syntax Description | <p>always (Optional) Always performs a soft inbound clear using stored updates, even if the neighbor supports the route refresh capability.</p> <p>inheritance-disable (Optional) Overrides configuration for this command that may be inherited from a neighbor group or address family group.</p> | |
| Command Default | Soft reconfiguration is not enabled. | |
| Command Modes | <p>IPv4 address family group configuration</p> <p>IPv6 address family group configuration</p> <p>IPv4 neighbor address family configuration</p> <p>IPv4 neighbor group address family configuration</p> <p>IPv6 neighbor group address family configuration</p> | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>To filter or modify some of the updates received from a neighbor, you configure an inbound policy using the route-policy (BGP) command. Configuring soft reconfiguration inbound causes the software to store the original unmodified route beside a route that is modified or filtered. This allows a “soft clear” to be performed after the inbound policy is changed. To perform a soft clear, use the clear bgp soft command with the in keyword specified. The unmodified routes are then passed through the new policy and installed in the BGP table.</p> | |
|  | | |
| Note | If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors. | |



Note The `bgp auto-policy-soft-reset` is enabled by default. A soft clear is done automatically when the inbound policy configured with the **route-policy** (BGP) command is changed. This behavior can be changed by disabling the auto-policy-soft-reset using the **bgp auto-policy-soft-reset disable** command.

If the neighbor supports the route refresh capability, then the original routes are not stored because they can be retrieved from the neighbor through a route refresh request. However, if the **always** keyword is specified, the original routes are stored even when the neighbor supports the route refresh capability.

If the **soft-reconfiguration inbound** command is not configured and the neighbor does not support the route refresh capability, then an inbound soft clear is not possible. In that case, the only way to rerun the inbound policy is to use the **clear bgp ip-address** command to reset the neighbor BGP session.



Note If there is an existing BGP session with a neighbor that does not support the route refresh capability, the session is terminated and a new one is initiated.



Note The extra routes stored as a result of configuring this command use more memory on the router.

If you configure this command for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID Operations

bgp read,
write

Examples

The following example shows inbound soft reconfiguration enabled for IP Version 4 (IPv4) unicast routes received from neighbor 10.108.1.1. The software stores all routes received in their unmodified form so that when an inbound soft clear is performed later, the stored information can then be used to generate a new set of modified routes.

```
RP/0/RP0/CPU0:router(config)# router bgp 100
RP/0/RP0/CPU0:router(config-bgp)# neighbor 10.108.1.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

| Command | Description |
|--------------------------------------|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |

| Command | Description |
|--|---|
| bgp auto-policy-soft-reset disable, on page 29 | Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified. |
| clear bgp, on page 82 | Resets a BGP connection using a soft or hard reset. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| route-policy (BGP), on page 178 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |

speaker-id

To allocate a speaker process to a neighbor, use the **speaker-id** command in the appropriate configuration mode. To remove the speaker process from a neighbor, use the **no** form of this command.

speaker-id *id*
no speaker-id [*id*]

| | |
|---------------------------|--|
| Syntax Description | <i>id</i> ID of the speaker process. Range is 1 to 15. |
|---------------------------|--|

| | |
|------------------------|---------------|
| Command Default | Default is 0. |
|------------------------|---------------|

| | |
|----------------------|-----------------------------|
| Command Modes | Session group configuration |
|----------------------|-----------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

| | |
|-----------------|--|
| Examples | The following example shows how to allocate speaker process 3 to neighbor 192.168.40.24: |
|-----------------|--|

```
RP/0/RP0/CPU0:router(config)# router bgp 109  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24  
RP/0/RP0/CPU0:router(config-bgp-nbr)# speaker-id 3
```

table-policy

To apply a routing policy to routes being installed into the routing table, use the **table-policy** command in an appropriate configuration mode. To disable applying a routing policy when installing routes into the routing table, use the **no** form of this command.

table-policy *policy-name*
no table-policy [*policy-name*]

| Syntax Description | <i>policy-name</i> Name of the routing policy to apply. | | | | |
|---------------------------|--|---------|--------------|---------------|------------------------------|
| Command Default | No policy is applied when routes are installed into the routing table. | | | | |
| Command Modes | IPv4 address family configuration IPv6 address family configuration | | | | |
| Command History | <table> <tr> <th>Release</th><th>Modification</th></tr> <tr> <td>Release 5.0.1</td><td>This command was introduced.</td></tr> </table> | Release | Modification | Release 5.0.1 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.1 | This command was introduced. | | | | |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

Table policy provides users with the ability to drop routes from the RIB based on match criteria. This feature can be useful in certain applications and should be used with caution as it can easily create a routing ‘black hole’ where BGP advertises routes to neighbors that BGP does not install in its global routing table and forwarding table.

Use the **table-policy** command to modify route attributes as the routes are installed into the routing table by Border Gateway Protocol (BGP). Commonly, it is used to set the traffic index attribute.

| Task ID | Task ID | Operations |
|---------|---------|----------------|
| | bgp | read, write |

Examples

The following example shows how to apply the set-traffic-index policy to IPv4 unicast routes being installed into the routing table:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)# table-policy set-traffic-index
```

Related Commands

| Command | Description |
|---------------------------|--|
| route-policy (RPL) | Defines a route policy and enters route policy configuration mode. |

timers (BGP)

To set the timers for a specific Border Gateway Protocol (BGP) neighbor, use the **timers** command in an appropriate configuration mode. To set the timers to the default values, use the **no** form of this command.

timers *keepalive hold-time*
no timers [*keepalive hold-time*]

Syntax Description

keepalive Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 65535.

hold-time Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 65535.

Command Default

keepalive : 60 seconds

hold-time : 180 seconds

Use the **timers bgp** command to override the default values.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbours that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

Task Operations

bgp read,
write

Examples

The following example shows how to change the keepalive timer to 70 seconds and the hold-time timer to 210 seconds for the BGP peer 192.168.40.24:

```
RP/0/RP0/CPU0:router(config)# router bgp 109
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# timers 70 210
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| timers bgp, on page 353 | Adjusts BGP network timers for all BGP neighbors. |

timers bgp

To change the default timer values for Border Gateway Protocol (BGP) neighbors, use the **timers bgp** command in an appropriate configuration mode. To set the default timers to the default values, use the **no** form of this command.

timers bgp *keepalive hold-time*
no timers bgp [*keepalive hold-time*]

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | <p><i>keepalive</i> Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 65535.</p> <p><i>hold-time</i> Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 65535.</p> | |
| Command Default | <p><i>keepalive</i> : 60 seconds</p> <p><i>hold-time</i> : 180 seconds</p> | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the timers bgp command to adjust the default timer times used by all BGP neighbors. The values can be overridden on particular neighbors using the timers command in the neighbor configuration mode.</p> <p>The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.</p> <p>The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.</p> <p>In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbors that will be using sessions with the non-default timers.</p> <p>Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to</p> | |

CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

Task ID

bgp read,
write

Examples

The following example shows how to configure a default keepalive time of 30 seconds and a default hold time of 90 seconds:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# timers bgp 30 90
```

Related Commands

| Command | Description |
|---|--|
| timers (BGP), on page 351 | Adjusts BGP network timers for a BGP neighbor. |

ttl-security

To configure a router to check the time-to-live (TTL) field in incoming IP packets for the specified external Border Gateway Protocol (eBGP) peer, use the **ttl-security** command in an appropriate configuration mode. To disable TTL verification, use the **no** form of this command.

ttl-security [**inheritance-disable**]
no ttl-security [**inheritance-disable**]

| Syntax Description | inheritance-disable (Optional) Prevents the ttl-security command from being inherited from a session group or neighbor group. | | | | |
|---------------------------|--|---------|--------------|---------------|------------------------------|
| Command Default | TTL verification is not enabled for eBGP peers. | | | | |
| Command Modes | Neighbor configuration Neighbor group configuration Session group configuration | | | | |
| Command History | <table><tr><th>Release</th><th>Modification</th></tr><tr><td>Release 5.0.0</td><td>This command was introduced.</td></tr></table> | Release | Modification | Release 5.0.0 | This command was introduced. |
| Release | Modification | | | | |
| Release 5.0.0 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the ttl-security command to enable a lightweight security mechanism to protect eBGP peering sessions from CPU utilization-based and other resource exhaustion-based attacks. These types of attacks are typically brute-force Denial of Service (DoS) attacks that attempt to disable the network by flooding devices in the network with IP packets that contain forged source and destination IP addresses in the packet headers.</p> <p>This command leverages existing behavior in IP packets. For a given IP packet, the TTL count of the packet always is equal to or less than the TTL count when the packet originated, a behavior that is considered impossible to circumvent. Therefore, a packet received with a TTL count equal to the maximum TTL value of 255 can be sent only by a directly adjacent peer. When the ttl-security command is configured for an eBGP neighbor that is directly adjacent, the router accepts only IP packets with a TTL count that is equal to the maximum TTL value.</p> <p>The ttl-security command secures the eBGP session in the incoming direction only. In the outbound direction, it causes packets to be sent only with the maximum TTL value so that the BGP neighbor can also verify the TTL value of incoming packets. When this command is enabled, BGP establishes or maintains a session only if the TTL value in the IP packet header is equal to the maximum TTL value. If the value is less than the maximum TTL value, the packet is discarded and an Internet Control Message Protocol (ICMP) message is not generated. This behavior is designed because a response to a forged packet is not necessary.</p> | | | | |



Note The **ttl-security** command must be configured on each participating router. Failure to configure this command on both ends of the BGP session results in the session progressing as far as the OpenSent or OpenConfirm state, remaining there until the hold time expires.

The following restrictions apply to the configuration of this command:

- The **ttl-security** command should not be configured for a peer that is already configured with the **neighbor ebgp-multihop** command. The simultaneous configuration of these commands is permitted; however, the **ttl-security** command overrides the **ebgp-multihop** command.
- This command is not supported for internal BGP (iBGP) peers.
- This command is not effective against attacks from a directly adjacent peer that has been compromised.

If you configure this command for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.



Note If the **ttl-security** command is configured on a neighbor to which the router has an established connection or the router is in the process of establishing a connection, the session must be cleared using the **clear bgp** command.

Task ID

Task ID Operations

bgp read,
write

Examples

The following example shows how to enable TTL security for eBGP neighbor 192.168.223.7:

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.7
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65507
RP/0/RP0/CPU0:router(config-bgp-nbr)# ttl-security
```

The following example shows how to enable TTL security for multiple eBGP neighbors using a session group:

```
RP/0/RP0/CPU0:router(config)# router bgp 65534
RP/0/RP0/CPU0:router(config-bgp)# session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-sngrp)# ttl-security
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.1
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65501
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.2
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65502
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 192.168.223.3
```

```
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 65503
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

| Command | Description |
|---|--|
| ebgp-multihop, on page 116 | Accepts and attempts BGP connections to external peers residing on networks that are not directly connected. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| show lpts flows | Displays information about locally terminated packet flows, including the minimum TTL value expected. |

update limit

To set upper bound on transient memory usage for update generation, use the **update limit** command in router configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit *update-limit-MB*
no update limit

| | |
|---------------------------|---|
| Syntax Description | <i>update-limit-MB</i> Sets the update limit in megabytes (MB). Range is 16 to 2048 MB. |
|---------------------------|---|

| | |
|------------------------|---------------------------------|
| Command Default | Default update limit is 512 MB. |
|------------------------|---------------------------------|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

Use the **update limit** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

| | | |
|----------------|----------------|------------------|
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to set the update limit as 1024 MB:

```
RP/0/RP0/CPU0:router(config)# router bgp 65000
RP/0/RP0/CPU0:router(config-bgp)#update limit 1024
```

| | | |
|-------------------------|--|---|
| Related Commands | Command | Description |
| | update limit address-family, on page 359 | Sets upper bound on transient memory usage for update generation for an address family. |
| | update limit sub-group, on page 360 | Sets upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups. |

update limit address-family

To set upper bound on transient memory usage for update generation for an address family, use the **update limit address-family** command in an appropriate address-family configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit address-family *update-limit-MB*
no update limit address-family

| | | |
|---------------------------|---|------------------------------|
| Syntax Description | <i>update-limit-MB</i> Sets the update limit in megabytes (MB). Range is 4 MB to 2048 MB. | |
| Command Default | Default update limit is 256 MB. | |
| Command Modes | IPv4 address family configuration IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the update limit address-family command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.</p> | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to set the update limit as *512 MB* for address family IPv4 unicast:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-af)#update limit address-family 512
```

| | | |
|-------------------------|---|---|
| Related Commands | Command | Description |
| | update limit, on page 358 | Sets upper bound on transient memory usage for update generation. |
| | update limit sub-group, on page 360 | Sets upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups. |

update limit sub-group

To set upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups, use the **update limit sub-group** command in an appropriate address-family configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit sub-group {ebgp | ibgp} *update-limit-MB*
no update limit sub-group {ebgp | ibgp}

| | | |
|---------------------------|--|---|
| Syntax Description | ebgp | Specifies the update limit for eBGP sub-groups. |
| | ibgp | Specifies the update limit for iBGP sub-groups. |
| | <i>update-limit-MB</i> | Sets the update limit in megabytes (MB). Range is 1 MB to 512 MB. |
| Command Default | Default update limit is 32 MB. | |
| Command Modes | IPv4 address family configuration | |
| | IPv6 address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | Use the update limit sub-group command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to set the update limit as *256 MB* for eBGP sub-group under address family IPV4 unicast:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#address-family ipv4 unicast
```



```
RP/0/RP0/CPU0:router(config-bgp-af)#update limit sub-group ebgp 256
```

Related Commands

| Command | Description |
|--|---|
| update limit, on page 358 | Sets upper bound on transient memory usage for update generation. |
| update limit address-family, on page 359 | Sets upper bound on transient memory usage for update generation for an address family. |

update in error-handling basic disable

To disable inbound update message basic error handling for eBGP or iBGP neighbors, use the **update in error-handling basis disable** command in router configuration mode. To enable inbound update message basic error handling, use the **no** form of this command.

update in error-handling basic {ebgp | ibgp} disable
no update in error-handling basic {ebgp | ibgp} disable

| | | |
|---------------------------|---|---|
| Syntax Description | ebgp | Specifies inbound update message basic error handling for eBGP neighbors. |
| | ibgp | Specifies inbound update message basic error handling for iBGP neighbors. |
| Command Default | Inbound update message basic error handling is enabled. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to disable inbound update message basic error handling for eBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling basic ebgp disable
```

This example shows how to disable inbound update message basic error handling for iBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling basic ibgp disable
```

update in error-handling extended

To enable inbound update message extended error handling for eBGP or iBGP neighbors, use the **update in error-handling extended** command in router configuration mode. To disable inbound update message error handling, use the **no** form of this command.

update in error-handling extended {ebgp | ibgp}
no update in error-handling extended {ebgp | ibgp}

| | | |
|--------------------|---|--|
| Syntax Description | ebgp | Specifies to enable inbound update message extended error handling for eBGP neighbors. |
| | ibgp | specifies to enable inbound update message extended error handling for iBGP neighbors. |
| Command Default | Inbound update message extended error handling is disabled. | |
| Command Modes | Router configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

This example shows how to enable inbound update message extended error handling for eBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling extended ebgp
```

This example shows how to enable inbound update message extended error handling for iBGP neighbors:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update in error-handling extended ibgp
```

update in labeled-unicast equivalent

To treat SAFIs unicast and labelled unicast for a given neighbor as same, use the **update in labeled-unicast equivalent** command in neighbor configuration mode. This command was introduced for interoperability between Cisco IOS XR and IOS XE systems. For example, If the XE router advertises ipv4 unicast and then withdraw with the ipv4 label-unicast of the same route, with knob enabled, then the BGP in the receiving XR router treats both of them same. Hence BGP removes the prefix information from its database.

update in labeled-unicast equivalent [inheritance-disable]

| | | |
|---------------------------|---|---|
| Syntax Description | equivalent | IPv4/IPv6 labeled-unicast inbound updates (paths) treated equivalent to unicast updates (paths) |
| | inheritance-disable | Does not inherit this unicast configuration from parent group |
| Command Default | None | |
| Command Modes | Neighbor configuration mode | |
| | Neighbor group configuration mode | |
| Command History | Release | Modification |
| | Release 7.0.1 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| Task ID | Task ID | Operation |
| | bgp | read, write |

Example

The following example shows how to treat IPv4 labeled-unicast inbound updates equivalent to unicast updates (paths):

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RP0/CPU0:router(config-bgp-nbr)#update in labeled-unicast equivalentexit
```

update out logging

To enable logging of update generation events, use the **update out logging** command in router configuration mode. To disable the logging of update generation events, use the **no** form of this command.

update out logging
no update out logging

| | |
|---------------------------|--|
| Syntax Description | This command has no keywords or arguments. |
|---------------------------|--|

| | |
|------------------------|--|
| Command Default | Update generation event logging is disabled. |
|------------------------|--|

| | |
|----------------------|----------------------|
| Command Modes | Router configuration |
|----------------------|----------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.0.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

| Task ID | Task ID | Operation |
|---------|---------|----------------|
| | bgp | read, write |

This example shows how to enable logging of update generation events:

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#router bgp 100
RP/0/RP0/CPU0:router(config-bgp)#update out logging
```

update-source

To allow internal Border Gateway Protocol (iBGP) sessions to use the primary IP address from a particular interface as the local address when forming an iBGP session with a neighbor, use the **update-source** command in an appropriate configuration mode. To set the chosen local IP address to the nearest interface to the neighbor, use the **no** form of this command.

update-source *type interface-path-id*
no update-source [*type interface-path-id*]

Syntax Description

| | |
|--------------------------|--|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. |
| Note | Use the show interfaces command to see a list of all interfaces currently configured on the router. |
| | For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

Best local address

Command Modes

Neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **update-source** command is commonly used with the loopback interface feature for iBGP sessions. The loopback interface is defined, and the interface address is used as the endpoint for a BGP session through the **update-source** command. This mechanism allows a BGP session to remain up even if the outbound interface goes down, provided there is another route to the neighbor.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|----------------|
| bgp | read, write |

Examples

The following example shows how to configure this router to use the IP address from the Loopback0 interface when trying to open a session with neighbor 172.20.16.6:

```
RP/0/RP0/CPU0:router(config)# router bgp 110  
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6  
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 110  
RP/0/RP0/CPU0:router(config-bgp-nbr)# update-source Loopback0
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |

use

To inherit configuration from a neighbor group, session group, or address family group, use the **use** command in an appropriate configuration mode. To discontinue inheritance from a group, use the **no** form of this command.

```
use {af-group group-name | neighbor-group group-name | session-group group-name e}
no use {af-group [group-name] | neighbor-group [group-name] | session-group [group-name]}
```

| | | |
|---------------------------|---|--|
| Syntax Description | af-group | Specifies an address family group. |
| | <i>group-name</i> | Name of the neighbor group, session group, or address family group from which you want to inherit configuration. |
| | neighbor-group | Specifies a neighbor group. |
| | session-group | Specifies a session group. |
| Command Default | Inheritance of group characteristics does not occur. | |
| Command Modes | For use af-group version: | |
| | Address family group configuration | |
| | Neighbor address family configuration | |
| | Neighbor group address family configuration | |
| | For use neighbor-group version: | |
| | Neighbor group configuration | |
| | Neighbor configuration | |
| | For use session-group version: | |
| | Neighbor group configuration | |
| | Neighbor configuration | |
| | Session-group configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. | |
| | The use command configures inheritance of configuration from an address family group, neighbor group, or session group, which means that any configuration for the group also takes effect for the user of the group. | |

The configuration inherited depends on the type of group that is specified. The group types are described in the following sections:

Address Family Group

An address family group can specify a configuration for only a single address family. The address family specified when the address family group was defined (through the **af-group** command) must match the address family from which the group is used.

Neighbor Group

A neighbor group (like a neighbor) can have address family-independent configuration and address family-specific configuration. All of these configurations could be inherited.

Session Group

A session group can have only address family-independent configuration and thus only address family-independent configuration is inherited from it.

The following rules govern inheritance to resolve possible conflicting configuration:

1. If a command is configured directly on the neighbor that is using group configuration, the command overrides the value that would be normally inherited from the group.
2. If the neighbor is configured to use a session group (for address family-independent configuration) or an address family group (for address family-specific configuration) and the command is configured for the session group or address family group, that configuration is used.
3. The neighbor group configuration is used:
 - If the command is not configured directly on the neighbor and the neighbor is not using a session group (for address family-independent configuration) or an af-group (for address family-specific configuration).
 - The neighbor is using a neighbor group and the command is configured on the neighbor group.

Typically, all configuration for a neighbor group is inherited, but some characteristics may be masked by a session group or address family group. For an example of this configuration, see the “Examples” section.

If the neighbor is using both a session group and a neighbor group and a specific command is configured for the neighbor group but not for the session group, then the configuration for the neighbor group does not take effect. The session group “hides” all address family-independent configuration on the neighbor group and prevents it from being inherited. Similarly, the use of an address family group hides any address family-specific configuration that may otherwise be inherited from a neighbor group for that address family.

In addition to neighbors using groups, it is possible to build a hierarchy by having groups use other groups. The following hierarchical groups are permitted:

- Session groups may use other session groups.
- Address family groups may use other address family groups.
- Neighbor groups may use other neighbor groups.
- Neighbor groups may use session groups and address family groups.



Note

Within the Cisco IOS XR system configuration architecture, do not combine the **remote-as** command and the **no use neighbor-group** command in the same commit, or the **remote-as** command and the **no use session-group** command in the same commit.

Task ID**Task ID Operations**

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to define a session group session1 and configure neighbor 172.168.40.24 to use session1. As a result, the session1 configuration takes effect on the neighbor also.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RP0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

The following example is similar to the previous example, but in this case the **timers** command on the session group does not take effect on the neighbor because it is overridden by a **timers** command directly configured for the neighbor.

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RP0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/RP0/CPU0:router(config-bgp-nbr)# timers 60 180
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

The following example shows a neighbor inheriting configuration from a session group that likewise inherits configuration from another session group. The configuration from both session groups take effect on the neighbor:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# session-group session2
RP/0/RP0/CPU0:router(config-bgp-sngrp)# use session-group session1
RP/0/RP0/CPU0:router(config-bgp-sngrp)# update-source Loopback0
RP/0/RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# use session-group session2
RP/0/RP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands


| Command | Description |
|--|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| remote-as (BGP), on page 173 | Creates a BGP neighbor and begins the exchange of routing information. |
| show bgp af-group, on page 220 | Displays information about BGP configuration for address family groups. |
| show bgp neighbor-group, on page 252 | Displays information about the BGP configuration for neighbor groups. |
| show bgp neighbors, on page 256 | Displays information about BGP neighbors. |
| show bgp session-group, on page 316 | Displays information about the BGP configuration for session groups. |

weight

To assign a weight to routes received from a neighbor, use the **weight** command in an appropriate configuration mode. To remove the **weight** command from the configuration file and restore the system to its default condition in which the software assigns the default weight to routes, use the **no** form of this command.

weight *weight-value*

no weight [*weight-value*]

| | | |
|---|---|------------------------------|
| Syntax Description | <i>weight-value</i> Weight to assign. Range is 0 to 65535. | |
| Command Default | Routes learned through another Border Gateway Protocol (BGP) peer have a default weight of 0 and routes sourced by the local router have a default weight of 32768. | |
| Command Modes | IPv4 address family group configuration IPv6 address family group configuration IPv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration | |
| Command History | Release | Modification |
| | Release 5.0.0 | This command was introduced. |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>The weight of a route is a Cisco-specific attribute. It is used in the best-path selection process (as the strongest tie-breaker). See the <i>Implementing BGP on</i> module of the <i>Routing Configuration Guide for Cisco NCS 6000 Series Routers</i> for information on best path. If there are two BGP routes with the same network layer reachability information (NLRI), the route with the higher weight is always chosen no matter what the value of other BGP attributes. Weight only has significance on the local router. Weight is assigned locally to the router, is a value that only makes sense to the specific router, is not propagated or carried through any route updates, and never is sent between BGP peers (even within the same AS).</p> | |
|  | | |
| Note | <p>If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors.</p> <p>The weight assigned to individual routes can be further manipulated in the inbound route policy of a neighbor using the set weight command. The set weight command sets the weight directly. If you have particular neighbors that you want to prefer for most of your outbound traffic, you can assign a higher weight to all routes learned from that neighbor.</p> <p>The weight assigned to individual routes may be modified by using an inbound routing policy.</p> | |



Note For weight changes to take effect, you may need to use the [clear bgp soft, on page 101](#) command.

If this command configures a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

Task ID Operations

| | |
|-----|----------------|
| bgp | read, write |
|-----|----------------|

Examples

The following example shows how to assign a weight of 50 to all IP Version 4 (IPv4) unicast routes learned through 172.20.16.6:

```
RP/0/RP0/CPU0:router(config)# router bgp 1
RP/0/RP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# weight 50
RP/0/RP0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 10 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| clear bgp, on page 82 | Resets a group of BGP neighbors. |
| neighbor-group, on page 141 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 189 | Creates a session group and enters session group configuration mode. |
| set weight | Sets the weight for BGP routes. |

 weight