



BGP Route Reflector Commands

This chapter provides details of the commands used for configuring Border Gateway Protocol (BGP) Route Reflector (RR).

- [address-family \(BGP\), on page 2](#)
- [additional-paths selection, on page 5](#)
- [keychain, on page 6](#)
- [neighbor \(BGP\), on page 8](#)
- [remote-as \(BGP\), on page 9](#)
- [route-reflector-client, on page 11](#)
- [router bgp, on page 13](#)
- [show bgp advertised, on page 14](#)
- [show bgp neighbors, on page 20](#)
- [show bgp paths, on page 35](#)
- [show bgp policy, on page 37](#)
- [show bgp route-policy, on page 44](#)
- [show bgp summary, on page 48](#)
- [table-policy, on page 52](#)
- [update-source, on page 53](#)
- [next-hop-self, on page 55](#)

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 unicast | ipv4 multicast | ipv4 labeled-unicast | ipv4 tunnel | vpnv4 unicast
}
no address-family { ipv4 unicast | ipv4 multicast | ipv4 labeled-unicast | ipv4 tunnel | vpnv4
unicast }
```

Syntax Description		
	ipv4 unicast	Specifies IP Version 4 (IPv4) unicast address prefixes.
	ipv4 multicast	Specifies IPv4 multicast address prefixes.
	ipv4 labeled-unicast	Specifies IPv4 labeled-unicast address prefixes. This option is available in IPv4 neighbor configuration mode and VRF neighbor configuration mode.
	ipv4 tunnel	Specifies IPv4 tunnel address prefixes.
	vpn4 unicast	Specifies VPN Version 4 (VPNv4) unicast address prefixes. This option is not available in VRF or VRF 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
- VRF configuration
- VRF neighbor configuration (IPv4 address families)

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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.

Table 1: Address Family Submode Support

Address Family	Supported in Router Submode	Supported in Neighbor Submode	Comments
ipv4 unicast	Yes	Yes	-
ipv4 multicast	Yes	Yes	-
ipv4 tunnel	Yes	Yes	-
ipv4 labeled-unicast	Yes	Yes	The ipv4 labeled-unicast address family can be configured only as a neighbor address family; however, the ipv4 unicast address family must be configured as the router address family first.
vpn4 unicast	Yes	Yes	-

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, multicast, labeled-unicast, and VPNv4 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:hostname# configure
RP/0/RP0:hostname(config)# router bgp 100
RP/0/RP0:hostname(config-bgp)# address-family ipv4 unicast
RP/0/RP0:hostname(config-bgp-af)#
```

The following example shows how to activate IPv4 multicast for neighbor 10.0.0.1 and place the router in neighbor address family configuration mode for the IPv4 multicast address family:

```
RP/0/RP0:hostname# configure
RP/0/RP0:hostname# router bgp 1
RP/0/RP0:hostname(config-bgp)# address-family ipv4 multicast
RP/0/RP0:hostname(config-bgp-af)# exit
RP/0/RP0:hostname(config-bgp)# neighbor 10.0.0.1
RP/0/RP0:hostname(config-bgp-nbr)# remote-as 1
RP/0/RP0:hostname(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RP0:hostname(config-bgp-nbr-af)#
```

additional-paths selection

To configure additional paths selection mode for a prefix, use the **additional-paths selection** command in the appropriate configuration mode. Use the **additional-paths selection** command with an appropriate route-policy to calculate backup paths and to enable Prefix Independent Convergence (PIC) functionality.

additional-paths selection route-policy *route-policy name*

Syntax Description	route-policy <i>route-policy name</i> Specifies the name of a route policy used for additional paths selection.				
Command Default	None				
Command Modes	IPv4 address family configuration				
Command History	<table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 6.5.25</td> <td>This command was introduced.</td> </tr> </tbody> </table>	Release	Modification	Release 6.5.25	This command was introduced.
Release	Modification				
Release 6.5.25	This command was introduced.				
Usage Guidelines	To configure additional paths selection mode for some or all prefixes, use the additional-paths selection command by specifying a route-policy.				

Example

The following example shows how to use the **additional-paths selection** command:

```
RP/0/RP0:hostname (config-bgp-af)# additional-paths selection route-policy a1
```

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	Release	Modification
	Release 6.1.42	This command was introduced.

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

Specify a key chain to enable key chain authentication between two BGP peers. Use the **keychain** command to implement hitless key rollover for authentication.

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.



Note	BGP only supports HMAC-MD5 and HMAC-SHA1-12 cryptographic algorithms.
-------------	---

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

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*

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

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
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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 submode, 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:hostname(config)# router bgp 1
RP/0/RP0:hostname(config-bgp)# session-group group2
RP/0/RP0:hostname(config-bgp-sngrp)# remote-as 1
RP/0/RP0:hostname(config-bgp-sngrp)# exit
RP/0/RP0:hostname(config-bgp)# neighbor 10.0.0.1
RP/0/RP0:hostname(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:hostname(config)# router bgp 1
RP/0/RP0:hostname(config-bgp)# session-group group2
RP/0/RP0:hostname(config-bgp-sngrp)# remote-as 1
RP/0/RP0:hostname(config-bgp-sngrp)# exit
RP/0/RP0:hostname(config-bgp)# neighbor 10.0.0.1
RP/0/RP0:hostname(config-bgp-nbr)# use session-group group2
```

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 [disable]
no route-reflector-client [disable]
```

Syntax Description	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 IPv4 neighbor address family configuration IPv4 neighbor group address family configuration VPNv4 address family group configuration VPNv4 neighbor address family configuration VPNv4 neighbor group address family configuration				
Command History	<table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 6.1.42</td> <td>This command was introduced.</td> </tr> </tbody> </table>	Release	Modification	Release 6.1.42	This command was introduced.
Release	Modification				
Release 6.1.42	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 the proper task IDs. If you suspect 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:hostname(config)# router bgp 140
RP/0/RP0:hostname(config-bgp)# neighbor 172.20.1.1
RP/0/RP0:hostname(config-bgp-nbr)# remote-as 140
RP/0/RP0:hostname(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0:hostname(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 group group1:

```
RP/0/RP0:hostname(config)# router bgp 140
RP/0/RP0:hostname(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RP0:hostname(config-bgp-afgrp)# route-reflector-client
RP/0/RP0:hostname(config-bgp-afgrp)# exit
RP/0/RP0:hostname(config-bgp)# neighbor 172.20.1.1
RP/0/RP0:hostname(config-bgp-nbr)# remote-as 140
RP/0/RP0:hostname(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RP0:hostname(config-bgp-nbr-af)# use af-group group1
RP/0/RP0:hostname(config-bgp-nbr-af)# route-reflector-client inheritance-disable
```

router bgp

To configure the Border Gateway Protocol (BGP) routing process, use the **router bgp** command in config mode. To remove all BGP configurations and terminate the BGP routing process, use the **no** form of this command.

```
router bgp as-number
no 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

Config

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines

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

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:hostname(config)# router bgp 120
```

show bgp advertised

To display advertisements for neighbors or a single neighbor, use the **show bgp advertised** command in config mode.

```
show bgp [ ipv4 { unicast | multicast | labeled-unicast | all | tunnel } | all { unicast | multicast
| all | labeled-unicast | tunnel } | vpnv4 unicast [ rd rd-address ] | vrf { vrf-name | all } [
ipv4 { unicast | labeled-unicast } ] [ rd rd-address ] ] advertised [ neighbor ip-address ] [
summary ]
```

Syntax Description	
ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
tunnel	(Optional) Specifies tunnel address prefixes.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd rd-address	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
<i>vrf-name</i>	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { unicast labeled-unicast }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
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.
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 Config

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines

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

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:hostname# 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 2: 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:hostname# 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 3: 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.

show bgp neighbors

To display information about Border Gateway Protocol (BGP) connections to neighbors, use the **show bgp neighbors** command in config mode.

```
show bgp [ ipv4 { unicast | multicast | labeled-unicast | all | tunnel } | all { unicast | multicast | all | labeled-unicast | tunnel } | vpnv4 unicast | vrf { vrf-name | all } [ ipv4 { unicast | labeled-unicast } ] ] neighbors [ performance-statistics | missing-eor ]
```

To show one neighbor:

```
show bgp [ ipv4 { unicast | multicast | labeled-unicast | all | tunnel } | all { unicast | multicast | all | labeled-unicast | tunnel } | vpnv4 unicast | vrf { vrf-name | all } [ ipv4 { unicast | labeled-unicast } ] ] neighbors ip-address [ advertised-routes | dampened-routes | flap-statistics | performance-statistics | received { prefix-filter | routes } | routes ]
```

To show default afi or safi neighbor:

```
show bgp neighbors ip-address [ configuration [ defaults ] [ nvgen ] | inheritance ]
```

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
<i>vrf-name</i>	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { unicast labeled-unicast }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast 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	Displays standby BGP information.

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

Config

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines



Note To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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.

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 ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp neighbors** command:

```
RP/0/RP0:hostname# 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
  Address family IPv4 Multicast: 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%

For Address Family: IPv4 Multicast
BGP neighbor version 23
Update group: 1.2
Route refresh request: received 0, sent 0
Policy for incoming advertisements is pass-all
Policy for outgoing advertisements is pass-all
2 accepted prefixes, 2 are bestpaths
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 131072
Threshold for warning message 75%

Connections established 9; dropped 8
Last reset 00:02:10, due to User clear requested (CEASE notification sent - administrative
reset)
Time since last notification sent to neighbor: 00:02:10
Error Code: administrative reset
Notification data sent:
  None
```

This table describes the significant fields shown in the display.

Table 4: 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.
local AS	<p>Autonomous system number of the local system.</p> <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) is 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	<p>TCP mode used in establishing the BGP session. The following valid TCP mode are supported:</p> <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.

Field	Description
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.
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. <p>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.</p>
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.

Field	Description
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.
AF-dependant capabilities	BGP capabilities that are specific to a particular address family. The following valid AF-dependent BGP capabilities are supported: <ul style="list-style-type: none"> • route refresh capability • route refresh capability OLD value
Outbound Route Filter	Neighbor has the Outbound Route Filter (ORF) capability for the specified address family. Details of the capabilities supported are also shown: 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. 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.
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.

Field	Description
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.
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	<p>Indicates whether the condition map selects routes that should be advertised, or routes that should not be advertised:</p> <p>Exist—Routes advertised if permitted by the condition route map.</p> <p>Non-exist—Routes advertised if denied by the condition route map.</p>
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	<p>Number of prefix updates that were suppressed because no transitive attributes changed from one best path to the next.</p> <p>Note Update suppression occurs only for external BGP neighbors.</p>
withdrawn	Number of prefixes withdrawn from the neighbor during the lifetime of the current connection with the neighbor.

Field	Description
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.
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:hostname# show bgp neighbors 10.0.101.75 advertised-routes

Network Next-hop From
10.10.0.0/8 10.0.101.1 10.0.101.1
10.11.0.0/8 10.0.101.3 10.0.101.3
10.12.0.0/8 10.0.101.5 10.0.101.5
```

The following is sample output from the show bgp neighbors command with the advertised-routes keyword:

```
RP/0/RP0:hostname# 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:hostname# 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 5: 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.

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.

The following is sample output from the **show bgp neighbors** command with the **dampened-routes** keyword:

```

RP/0/RP0:hostname# 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 6: 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.
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.

The following is sample output from the **show bgp neighbors** command with the **performance-statistics** keyword:

```
RP/0/RP0:hostname# 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 7: 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> <ul style="list-style-type: none"> S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> d—Path is dampened. 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>The third character may be:</p> <ul style="list-style-type: none"> i—Path was learned by an internal BGP (iBGP) session.
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> <ul style="list-style-type: none"> 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.

Field	Description
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 **configuration** keyword:

```
RP/0/RP0:hostname# 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 []
  address-family ipv4 multicast []
    policy pass-all in []
    policy pass-all out []
```

This table describes the significant fields shown in the display.

Table 8: 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.

show bgp paths

To display all the Border Gateway Protocol (BGP) paths in the database, use the **show bgp paths** command in config 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 Config

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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.

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:hostname# 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 9: 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 config mode.

```
show bgp [ ipv4 { unicast | multicast | labeled-unicast | all | tunnel } | all { unicast |
multicast | all | labeled-unicast | tunnel } | vpnv4 unicast [ rd rd-address ] | vrf { vrf-name
| all } [ ipv4 { unicast | labeled-unicast } ] [ rd rd-address ] ] policy [ neighbor ip-address ]
[ sent-advertisements | route-policy route-policy-name ] [ summary ]
```

Syntax	Description
ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd <i>rd-address</i>	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
<i>vrf-name</i>	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { unicast labeled-unicast }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
neighbor	(Optional) Previews advertisements for a single neighbor.
<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.
<i>route-policy-name</i>	(Optional) Name of the route policy.
summary	(Optional) Displays a summary of the BGP advertisements.

Command Default Advertisements for all neighbors are displayed if the **neighbor** *ip-address* 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 Config

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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.

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**

```
RP/0/RP0:hostname# 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 10: show bgp policy 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.
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:

```
RP/0/RP0:hostname# 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 11: 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.

Field	Description
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.
(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	<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.
Attributes after outbound policy was applied	<p>Displays attributes associated with the received route, after any outbound 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>

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 config mode.

```
show bgp [ ipv4 { unicast | multicast | labeled-unicast | all | tunnel }
| all { unicast | multicast | all | labeled-unicast | tunnel }
| vpnv4 unicast [ rd rd-address ] | vrf { vrf-name | all } [ ipv4 { unicast | labeled-unicast
} ] route-policy route-policy-name
```

Syntax Description	
ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
rd <i>rd-address</i>	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
<i>vrf-name</i>	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { unicast labeled-unicast }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families.
<i>route-policy-name</i>	Name of a route policy.

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 Config

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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.

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 :

```
RP/0/RP0:hostname# show bgp route-policy pl

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 12: show bgp route-policy Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.

Field	Description
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.
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.

Field	Description
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.

show bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the **show bgp summary** command in config mode.

```
show bgp [{ipv4{unicast | multicast | labeled-unicast | all | tunnel } | all {unicast | multicast | all |
labeled-unicast | tunnel } | vpnv4 unicast | vrf {vrf-name | all}][{ipv4 {unicast | labeled-unicast }}]
summary
```

Syntax Description	
ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress families, specifies prefixes for all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
<i>vrf-name</i>	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { unicast labeled-unicast }	(Optional) For VRF, specifies IPv4 unicast or labeled-unicast 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 Config

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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.

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 (RcvTblVer, 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 < RcvTblVer—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 Operations ID

bgp read

Examples

The following is sample output from the **show bgp summary** command:

```
RP/0/RP0:hostname#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 13: show bgp summary Field Descriptions

Field	Description
BGP router identifier	IP address of the router.

Field	Description
local AS number	Autonomous system number set by the router bgp, on page 13 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.
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.

Field	Description
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.
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>

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 VRF IPv4 address family configuration
----------------------	--

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines	To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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 an unwanted traffic drop 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:hostname(config)# router bgp 1
RP/0/RP0:hostname(config-bgp)# address-family ipv4 unicast
RP/0/RP0:hostname(config-bgp-af)# table-policy set-traffic-index
```

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.

Command Default Best local address

Command Modes Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History	Release	Modification
	Release 6.1.42	This command was introduced.

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect 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:hostname(config)# router bgp 110  
RP/0/RP0:hostname(config-bgp)# neighbor 172.20.16.6  
RP/0/RP0:hostname(config-bgp-nbr)# remote-as 110  
RP/0/RP0:hostname(config-bgp-nbr)# update-source Loopback0
```

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 VPNv4 address family group configuration IPv4 neighbor address family configuration VPNv4 neighbor address family configuration IPv4 neighbor group address family configuration IPv6 neighbor group address family configuration VPNv4 neighbor group address family configuration IPv4 labeled-unicast address family configuration IPv6 labeled-unicast address family configuration VRF labeled-unicast address family configuration
----------------------	--

Command History	<table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 6.5.33</td> <td>This command was introduced.</td> </tr> </tbody> </table>	Release	Modification	Release 6.5.33	This command was introduced.
Release	Modification				
Release 6.5.33	This command was introduced.				

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

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.

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.

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.

Configuring the **next-hop-self** command under IPv4 labeled-unicast, IPv6 labeled-unicast, or VRF labeled-unicast address family configuration mode enables next-hop-self for labeled prefixes advertised to an iBGP peer.

Task ID	Task ID	Operation
	bgp	read, write

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 192.0.2.2 to an address of the local router:

```
RP/0/RP0:hostname(config)#router bgp 1
RP/0/RP0:hostname(config-bgp)#neighbor 192.0.2.2
RP/0/RP0:hostname(config-bgp-nbr)#remote-as 1
RP/0/RP0:hostname(config-bgp)#bfd multiplier 3
RP/0/RP0:hostname(config-bgp)#bfd minimum-interval 2
RP/0/RP0:hostname(config-bgp-nbr)#bfd fast-detect
RP/0/RP0:hostname(config-bgp-nbr)#update-source hundredgige0/4/0/5.1
RP/0/RP0:hostname(config-bgp-nbr)#address-family ipv4 labeled unicast
RP/0/RP0:hostname(config-bgp-nbr)#route-reflector-client
RP/0/RP0:hostname(config-bgp-nbr)#next-hop-self
```