BGP Commands on Cisco IOS XR Software

This module describes the commands used to configure and monitor Border Gateway Protocol (BGP) for IP Version 4 (IPv4), IP Version 6 (IPv6), Virtual Private Network Version 4 (VPNv4), Virtual Private Network Version 6 (VPNv6), and multicast distribution tree (MDT) routing sessions.

For detailed information about BGP concepts, configuration tasks, and examples, see the *Implementing BGP on* Cisco IOS XR *Software* module in the *Cisco IOS XR Routing Configuration Guide for the Cisco XR 12000 Series Router*.

additional-paths install backup

Note   Effective with Release 4.0.0, the **additional-paths install backup** command was replaced by the **additional-paths selection** command. See the additional-paths selection, command for more information.

To install a backup path into the forwarding table and provide prefix independent convergence (PIC) in case of a PE-CE link failure, use the **additional-paths install backup** command in an appropriate address family configuration mode. To prevent installing the backup path, use the **no** form of this command. To disable prefix independent convergence, use the **disable** keyword.

additional-paths install backup [**disable**]

no additional-paths install backup

Syntax Description

disable	Disables installing backup path into the forwarding table.

Command Default   No default behavior or values

Command Modes   VRF IPv4 address family configuration

VRF IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.
Release 4.0.0	This command was replaced by the **additional-paths selection** command.

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.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable installing a backup path into the forwarding table in VPNv4 address family mode:

```
RP/0/0/CPU0:router#configure
RP/0/0/CPU0:router(config)#router bgp 100
RP/0/0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/0/CPU0:router(config-bgp-af)#additional-paths install backup
```

Related Commands

Command	Description
advertise best-external, page 18	Advertises the best-external path to the iBGP and route-reflector peers.
retain local-label, page 221	Retains the local label until the network is converged.

additional-paths receive

To configure receive capability of multiple paths for a prefix to the capable peers, use the **additional-paths receive** command in address-family configuration mode. To disable the receive capability, use the **no** form of this command. To disable add-path receive capability for all neighbors belonging to a particular VRF address-family, use the **disable** option.

additional-paths receive [disable]

no additional-paths receive

Syntax Description   This command has no arguments or keywords.

Command Default   No default behavior or values

Command Modes   IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 4.0.0	This command was introduced.

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

After enabling the receive capability, the session needs to be reset for the configuration to take into effect.

Task ID

Task ID	Operation
bgp	read, write

Examples   The following example shows how to enable additional paths receive capability under IPv4 unicast address family mode:

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:routerconfig)# router bgp 100
```

```
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# additional-paths receive
```

additional-paths selection

To configure additional paths selection capability for a prefix, use the **additional-paths selection** command in address-family configuration mode. To disable the additional-paths selection capability, use the **no** form of this command. To disable add-path selection for a particular VRF address-family, use the **disable** option.

additional-paths selection {**route-policy**_route-policy-name_| **disable**}

no additional-paths selection route-policy _route-policy-name_

Syntax Description

route-policy _route-policy-name_	Specifies the name of a route policy used for additional paths selection.
disable	Disables add-path selection for a particular VRF address-family.

Command Default

No default behavior or values

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 4.0.0	This command was introduced.

Usage Guidelines

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

To select the add-path mode for some or all prefixes, use the additional-paths selection command by specifying a route-policy.

Task ID

Task ID	Operation
bgp	read, write

Examples   The following example shows how to enable selection of additional paths:

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# additional-paths selection route-policy add-path-to-ibgp
```

additional-paths send

To configure send capability of multiple paths for a prefix to the capable peers, use the **additional-paths send** command in address-family configuration mode. To disable the send capability, use the **no** form of this command.

additional-paths send

no additional-paths send

Syntax Description   This command has no arguments or keywords.

Command Default   No default behavior or values

Command Modes   IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 4.0.0	This command was introduced.

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

After enabling the send capability, the session needs to be reset for the configuration to take into effect.

Task ID

Task ID	Operation
bgp	read, write

Examples   The following example shows how to enable additional paths send capability under IPv4 unicast address family mode:

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

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**| **ipv4mdt ipv6unicast ipv6multicast** | **ipv6 labeled- unicast**| **vpnv4** {**unicast**| **vpnv6**} **unicast**| **l2vpn vpls-vpws**}

no address-family {**ipv4 unicast**| **ipv4 multicast**| **ipv4 labeled-unicast**| **ipv4 tunnel**| **ipv4mdt ipv6unicast ipv6multicast** | **ipv6 labeled- unicast**| **vpnv4** {**unicast**| **vpnv6**} **unicast**| **l2vpn vpls-vpws**}

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.
ipv4 mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes. This option is available in router configuration mode and IPv4 neighbor configuration mode.
ipv6 unicast	Specifies IP Version 6 (IPv6) unicast address prefixes.
ipv6 multicast	Specifies IPv6 multicast address prefixes.
ipv6 labeled-unicast	Specifies IPv6 labeled-unicast address prefixes. This option is available in IPv6 neighbor configuration mode.
vpnv4 unicast	Specifies VPN Version 4 (VPNv4) unicast address prefixes. This option is not available in VRF or VRF neighbor configuration mode.
vpnv6 unicast	Specifies VPN Version 6 (VPNv6) unicast address prefixes. This option is not available in VRF or VRF neighbor configuration mode.
l2vpn vpls-vpws	Specifies L2VPN vpls-vpws address prefixes.

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 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF and VRF neighbor configuration modes. The **vpnv4 unicast** keywords were added.
Release 3.4.0	The **labeled-unicast** keyword was supported
Release 3.5.0	The following keywords were added: • **vpnv6 unicast** • **ipv4 tunnel** • **ipv4 mdt** • **ipv6 labeled-unicast**
Release 3.7.0	The Address Family Submode Support table was added.
Release 3.9.0	L2VPN Address Family support was added.

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

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	—

Address Family	Supported in Router Submode	Supported in Neighbor Submode	Comments
ipv4 mdt	yes	yes	—
ipv4 tunnel	yes	yes	—
ipv4 labeled-unicast	no	yes	The ipv4 labeled-unicast address family can be configured only as a neighbor address family; however, it requires that the ipv4 unicast address family be configured as the router address family first.
vpnv4 unicast	yes	yes	—
ipv6 unicast	yes	yes	—
ipv6 multicast	yes	yes	—
ipv6 labeled-unicast	no	yes	The ipv6 labeled-unicast address family can be configured only as a neighbor address family; however, it requires that the ipv6 unicast address family be configured as the router address family first.
vpnv6 unicast	yes	yes	—
l2vpn vpls-vpws	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. IPv6 neighbor sessions support IPv6 unicast and multicast 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/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(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/0/0/CPU0:router# configure
RP/0/0/CPU0:router# router bgp 1
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-af)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbr-af)#
```

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

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 12
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 tunnel
RP/0/0/CPU0:router(config-bgp-af)#
```

advertise best-external

To advertise the best–external path to the iBGP and route-reflector peers, when a locally selected bestpath is from an internal peer, use the **advertise best-external** command in an appropriate address family configuration mode. To prevent advertising the best–external path, use the **no** form of this command. To disable advertising the best–external path, use the **disable** keyword.

advertise best-external [disable]

no advertise best-external

Syntax Description

disable	Disables best–external configuration for the VRF.

Command Default

No default behavior or values

Command Modes

VRF IPv4 address family configuration

VRF IPv6 address family configuration

L2VPN address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

IPv4 address family configuration

IPv6 address family configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.
Release 4.0.0	This command was supported in global IPv4 and IPv6 unicast address-families.

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.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable advertising the best–external path VPNv4 unicast address family mode:

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# advertise best-external
```

Related Commands

Command	Description
additional-paths install backup, page 7	Installs a backup path into the forwarding table and provides prefix independent convergence (PIC) in case of a PE-CE link failure.
retain local-label, page 221	Retains the local label until the network is converged.

advertisement-interval

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

advertisement-interval *seconds*

no advertisement-interval [*seconds*]

Syntax Description

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

Command Default

Default minimum interval:

For internal BGP (iBGP) peers is 0 seconds

For external BGP (eBGP) peers is 30 seconds

For customer edge (CE) peers is 0 seconds

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

VRF neighbor configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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.

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

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router(config)# router bgp 5
RP/0/0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/0/CPU0:router(config-bgp-nbr)# advertisement-interval 10
```

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.

af-group

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

af-group *af-group-name* **address-family**

no af-group

Syntax Description

af-group-name	Address family group name.
address-family	Enters address family configuration mode.
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.
ipv4 tunnel	Specifies IPv4 tunnel address prefixes.
ipv4 mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6 unicast	Specifies IP Version 6 (IPv6) unicast address prefixes.
ipv6 multicast	Specifies IPv6 multicast address prefixes.
ipv6 labeled-unicast	Specifies IPv6 labeled unicast address prefixes.
vpnv4 unicast	Specifies VPN Version 4 (VPNv4) unicast address prefixes.
vpnv6 unicast	Specifies VPN Version 6 (VPNv6) unicast address prefixes.

Command Default   No BGP address family group is configured.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The **vpnv4 unicast** and **labeled-unicast** keywords were added.
Release 3.4.0	The **labeled-unicast** keyword was supported.

Release	Modification
Release 3.5.0	The **vpnv6 unicast**, **ipv6 labeled-unicast**, **ipv4 tunnel**, and **ipv4 mdt** keywords were added.

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
neighbor (BGP), page 180	Enters neighbor configuration mode for configuring BGP routing sessions.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.
use, page 455	Inherits configuration from a neighbor group, session group, or address family group.

aggregate-address

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

aggregate-address *address/mask-length* **[as-set] [as-confed-set] [summary-only] [route-policy** *route-policy-name***]**

no aggregate-address

Syntax Description

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

Command Default

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

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF IPv4 address family configuration mode.
Release 3.5.0	This command was supported in VRF IPv6 address family configuration mode.

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.

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

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

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

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

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

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
network (BGP), page 184	Specifies the list of networks for the BGP routing process.

Command	Description
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

allocate-label

To allocate Multiprotocol Label Switching (MPLS) labels for specific IPv4 unicast or IPv6 unicast or VPN routing and forwarding (VRF) IPv4 unicast routes so that the BGP router can send labels with BGP routes to a neighboring router configured for labeled-or VPN routing and forwarding (VRF) IPv6 unicast sessions, use the **allocate-label** command in the appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

allocate-label {**route-policy** *route-policy-name*| **all**}

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

Syntax Description

all	Allocates labels for all prefixes
route-policy *route-policy-name*	Uses a route policy to select prefixes for label allocation.

Command Default   No default behavior or values

Command Modes   IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.4.0	The **all** keyword was added.
	The command was supported in VRF IPv4 address family configuration mode.
Release 3.5.0	This command was supported in IPv6 address family configuration mode and VRF IPv6 address family configuration mode.

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 **allocate-label** command with a route policy to trigger BGP to allocate labels for all or a filtered set of global routes (as dictated by the route policy). The command enables autonomous system border routers (ASBRs) that have labeled unicast sessions to exchange Multiprotocol Label Switching (MPLS) labels with the routes to the other autonomous system (AS) in Layer 3 Virtual Private Network (L3VPN) inter-AS deployments.

Note The **allocate-label all** command is functionally equivalent to the **allocate-label route-policy** *route-policy-name* command when the route policy is a pass-all policy.

See *Cisco IOS XR MPLS Configuration Guide for the Cisco XR 12000 Series Router* for information on using the **allocate-label** command for L3VPN inter-AS deployments and carrier-supporting-carrier IPv4 BGP label distribution.

Task ID

Task ID	Operations
bgp	read, write

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

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

allowas-in

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

allowas-in [*as-occurrence-number*]

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

Syntax Description

as-occurrence-number	(Optional) Number of times a PE ASN is allowed. Range is 1 to 10.

Command Default

No default behavior or values

Command Modes

Address family group configuration

Neighbor address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.

Usage Guidelines

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

Hub and spoke VPN networks require looping back of routing information to the hub PE through the hub customer edge (CE). See *Cisco IOS XR MPLS Configuration Guide for the Cisco XR 12000 Series Router* for information on hub and spoke VPN networks. This looping back, in addition to the presence of the PE ASN, causes the looped-back information to be dropped by the hub PE.

The **allowas-in** command prevents the looped-back information from being dropped by replacing the neighbor autonomous system number (ASN) with the PE ASN in the AS path. This allows the VPN customer to see a specified number of occurrences of the PE ASN in the AS path.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to allow five occurrences of the PE ASN:

```
RP/0/0/CPU0:router(config)# router bgp 6
RP/0/0/CPU0:router(config-bgp)# af-group group_1 address-family vpnv4 unicast
```

```
RP/0/0/CPU0:router(config-bgp-afgrp)# allowas-in 5
```

as-format

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

as-format asdot

no as-format

Syntax Description

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

Command Default

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

Command Modes

Global configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.

Usage Guidelines

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

as-override

To configure a provider edge (PE) router to override the autonomous system number (ASN) of a site with the ASN of a provider, use the **as-override** command in VRF neighbor address family configuration mode. To restore the system to its default condition, use the **no** form of this command.

as-override [inheritance-disable]

no as-override [inheritance-disable]

Syntax Description

inheritance-disable	(Optional) Prevents the **as-override** command from being inherited from a parent group.

Command Default

Automatic override of the ASN is disabled.

Command Modes

VRF neighbor address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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 **as-override** command in conjunction with the site-of-origin (SoO) feature, identifying the site where a route originated, and preventing routing loops between routers within a VPN.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure an ASN override:

```
RP/0/0/CPU0:router(config)# router bgp 6
RP/0/0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# as-override
```

Related Commands

Command	Description
site-of-origin (BGP), page 433	Configures the site of origin filtering.

bfd (BGP)

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

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

bfd {**multiplier** | **minimum-interval**} *value*

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

Syntax Description

multiplier *value*	Specifies the BFD session's multiplier value for the neighbor.
minimum-interval *value*	Specifies the BFD session's minimum-interval value for the neighbor.

Command Default

No default per neighbor parameters are set.

Command Modes

Neighbor address family independent configuration

Command History

Release	Modification
Release 3.4.0	This command was introduced.
Release 3.6.0	The arguments **multiplier** and **minimum-interval** were added for the neighbor address family independent configuration.

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.

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

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

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

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

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

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

Examples   The following example shows how to specify the BFD session's multiplier value for the neighbor:

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 65000
RP/0/0/CPU0:router(config-bgp-nbrgrp)#neighbor 3.3.3.2
RP/0/0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 311
RP/0/0/CPU0:router(config-bgp-nbr)# bfd multiplier 7
RP/0/0/CPU0:router(config-bgp-nbr)# neighbor 5.5.5.2
RP/0/0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 318
RP/0/0/CPU0:router(config-bgp-nbr)# bfd multiplier 4
RP/0/0/CPU0:router(config-bgp-nbr)# vrf one
RP/0/0/CPU0:router(config-bgp-vrf)# neighbor 3.12.1.2
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# bfd minimum-interval 119
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# bfd multiplier 10
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# commit
```

```
  RP/0/0/CPU0:router# show bfd session
  Interface          Dest Addr         Local det time(int*mult)      State
                                          Echo          Async
  ------------------- --------------   ----------------  ---------------- ---------
  Gi0/2/0/2          3.3.3.2          2177ms(311ms*7)   14s(2s*7)       UP
  Gi0/2/0/2.1        3.12.1.2         1190ms(119ms*10)  20s(2s*10)      UP
  PO0/3/0/6          5.5.5.2          1272ms(318ms*4)   8s(2s*4)        UP
```

```
  RP/0/0/CPU0:router# show bfd session detail
  I/f: GigabitEthernet0/2/0/2, Location: 0/2/CPU0, dest: 3.3.3.2, src: 3.3.3.1
   State: UP for 0d:0h:4m:44s, number of times UP: 1
  Received parameters:
   Version: 1, desired tx interval: 2 s, required rx interval: 2 s
   Required echo rx interval: 1 ms, multiplier: 7, diag: None
   My discr: 524295, your discr: 524296, state UP, D/F/P/C/A: 0/0/0/1/0
  Transmitted parameters:
   Version: 1, desired tx interval: 2 s, required rx interval: 2 s
   Required echo rx interval: 1 ms, multiplier: 7, diag: None
   My discr: 524296, your discr: 524295, state UP, D/F/P/C/A: 0/0/0/1/0
  Timer Values:
   Local negotiated async tx interval: 2 s
   Remote negotiated async tx interval: 2 s
   Desired echo tx interval: 311 ms, local negotiated echo tx interval: 311 ms
   Echo detection time: 2177 ms(311 ms*7), async detection time: 14 s(2 s*7)
  Local Stats:
   Intervals between async packets:
     Tx: Number of intervals=100, min=1664 ms, max=2001 ms, avg=1838 ms
         Last packet transmitted 313 ms ago
     Rx: Number of intervals=100, min=1662 ms, max=2 s, avg=1828 ms
         Last packet received 1615 ms ago
   Intervals between echo packets:
     Tx: Number of intervals=100, min=181 ms, max=462 ms, avg=229 ms
         Last packet transmitted 289 ms ago
     Rx: Number of intervals=100, min=178 ms, max=461 ms, avg=229 ms
```

```
        Last packet received 287 ms ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 100, min=0 us, max=4 ms, avg=860 us
Session owner information:
  Client          Desired interval      Multiplier
  --------------- --------------------  --------------
  bgp-0           311 ms                7

I/f: GigabitEthernet0/2/0/2.1, Location: 0/2/CPU0, dest: 3.12.1.2, src: 3.12.1.1
 State: UP for 0d:0h:4m:44s, number of times UP: 1
Received parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 10, diag: None
 My discr: 524296, your discr: 524295, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 10, diag: None
 My discr: 524295, your discr: 524296, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
 Local negotiated async tx interval: 2 s
 Remote negotiated async tx interval: 2 s
 Desired echo tx interval: 119 ms, local negotiated echo tx interval: 119 ms
 Echo detection time: 1190 ms(119 ms*10), async detection time: 20 s(2 s*10)
Local Stats:
 Intervals between async packets:
   Tx: Number of intervals=100, min=1664 ms, max=2001 ms, avg=1838 ms
       Last packet transmitted 314 ms ago
   Rx: Number of intervals=100, min=1662 ms, max=2 s, avg=1828 ms
       Last packet received 1616 ms ago
 Intervals between echo packets:
   Tx: Number of intervals=100, min=120 ms, max=223 ms, avg=125 ms
       Last packet transmitted 112 ms ago
   Rx: Number of intervals=100, min=119 ms, max=223 ms, avg=125 ms
       Last packet received 110 ms ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 100, min=0 us, max=2 ms, avg=850 us
Session owner information:
  Client          Desired interval      Multiplier
  --------------- --------------------  --------------
  bgp-0           119 ms                10

I/f: GigabitEthernet0/3/0/6, Location: 0/3/CPU0, dest: 5.5.5.2, src: 5.5.5.1
 State: UP for 0d:0h:4m:50s, number of times UP: 1
Received parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 4, diag: None
 My discr: 786436, your discr: 786433, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
 Version: 1, desired tx interval: 2 s, required rx interval: 2 s
 Required echo rx interval: 1 ms, multiplier: 4, diag: None
 My discr: 786433, your discr: 786436, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
 Local negotiated async tx interval: 2 s
 Remote negotiated async tx interval: 2 s
 Desired echo tx interval: 318 ms, local negotiated echo tx interval: 318 ms
 Echo detection time: 1272 ms(318 ms*4), async detection time: 8 s(2 s*4)
Local Stats:
 Intervals between async packets:
   Tx: Number of intervals=100, min=1663 ms, max=2 s, avg=1821 ms
       Last packet transmitted 1740 ms ago
   Rx: Number of intervals=100, min=1663 ms, max=2001 ms, avg=1832 ms
       Last packet received 160 ms ago
 Intervals between echo packets:
   Tx: Number of intervals=100, min=181 ms, max=484 ms, avg=232 ms
       Last packet transmitted 44 ms ago
   Rx: Number of intervals=100, min=179 ms, max=484 ms, avg=232 ms
       Last packet received 41 ms ago
 Latency of echo packets (time between tx and rx):
   Number of packets: 100, min=0 us, max=3 ms, avg=540 us
Session owner information:
  Client          Desired interval      Multiplier
  --------------- --------------------  --------------
  bgp-0           318 ms                4
```

```
RP/0/0/CPU0:router# show bgp nei 3.3.3.2

BGP neighbor is 3.3.3.2
 Remote AS 500, local AS 65000, external link
 Remote router ID 16.0.0.1
  BGP state = Established, up for 00:05:01
  BFD enabled (session up): mininterval: 311 multiplier: 7
  Last read 00:00:56, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 8 messages, 0 notifications, 0 in queue
  Sent 9 messages, 1 notifications, 0 in queue
  Minimum time between advertisement runs is 30 seconds

 For Address Family: IPv4 Unicast
  BGP neighbor version 2
  Update group: 0.2
  AF-dependant capabilities:
    Graceful Restart Capability advertised and received
      Neighbor preserved the forwarding state during latest restart
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 360 seconds
      Remote Restart time is 120 seconds
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is pass-all
  1 accepted prefixes, 1 are bestpaths
  Prefix advertised 1, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%
  An EoR was not received during read-only mode

  Connections established 1; dropped 0
  Last reset 00:06:58, due to User clear requested (CEASE notification sent - administrative
reset)
  Time since last notification sent to neighbor: 00:06:58
  Error Code: administrative reset
  Notification data sent:
    None

 RP/0/0/CPU0:router# show bgp nei 5.5.5.2

 BGP neighbor is 5.5.5.2
 Remote AS 500, local AS 65000, external link
 Remote router ID 16.0.0.1
  BGP state = Established, up for 00:05:04
  BFD enabled (session up): mininterval: 318 multiplier: 4
  Last read 00:00:58, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 8 messages, 0 notifications, 0 in queue
  Sent 9 messages, 1 notifications, 0 in queue
  Minimum time between advertisement runs is 30 seconds

 For Address Family: IPv4 Unicast
  BGP neighbor version 2
  Update group: 0.2
  AF-dependant capabilities:
    Graceful Restart Capability advertised and received
      Neighbor preserved the forwarding state during latest restart
      Local restart time is 120, RIB purge time is 600 seconds
      Maximum stalepath time is 360 seconds
      Remote Restart time is 120 seconds
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is pass-all
  1 accepted prefixes, 0 are bestpaths
```

```
     Prefix advertised 1, suppressed 0, withdrawn 0, maximum limit 524288
     Threshold for warning message 75%
     An EoR was not received during read-only mode

     Connections established 1; dropped 0
   Last reset 00:07:01, due to User clear requested (CEASE notification sent - administrative
 reset)
     Time since last notification sent to neighbor: 00:07:01
     Error Code: administrative reset
     Notification data sent:
       None

 RP/0/0/CPU0:router# show bgp vrf one nei 3.12.1.2

 BGP neighbor is 3.12.1.2, vrf one
  Remote AS 500, local AS 65000, external link
  Remote router ID 16.0.0.1
    BGP state = Established, up for 00:05:06
    BFD enabled (session up): mininterval: 119 multiplier: 10
    Last read 00:00:01, hold time is 180, keepalive interval is 60 seconds
    Precedence: internet
    Neighbor capabilities:
      Route refresh: advertised and received
      4-byte AS: advertised and received
      Address family IPv4 Unicast: advertised and received
    Received 9 messages, 0 notifications, 0 in queue
    Sent 9 messages, 1 notifications, 0 in queue
    Minimum time between advertisement runs is 0 seconds

   For Address Family: IPv4 Unicast
    BGP neighbor version 2
    Update group: 0.2
    AF-dependant capabilities:
      Graceful Restart Capability advertised and received
        Neighbor preserved the forwarding state during latest restart
        Local restart time is 120, RIB purge time is 600 seconds
        Maximum stalepath time is 360 seconds
        Remote Restart time is 120 seconds
    Route refresh request: received 0, sent 0
    Policy for incoming advertisements is pass-all
    Policy for outgoing advertisements is pass-all
    1 accepted prefixes, 1 are bestpaths
    Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
    Threshold for warning message 75%
    An EoR was not received during read-only mode

    Connections established 1; dropped 0
   Last reset 00:07:04, due to User clear requested (CEASE notification sent - administrative
 reset)
     Time since last notification sent to neighbor: 00:07:04
     Error Code: administrative reset
     Notification data sent:
       None
```

bgp as-path-loopcheck

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

bgp as-path-loopcheck

no bgp as-path-loopcheck

Syntax Description   This command has no arguments or keywords.

Command Default   When you do not specify this command, loop checking is performed only for external peers.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to configure an autonomous system path for loop checking iBGP peers:

```
RP/0/0/CPU0:router(config)# router bgp 6
RP/0/0/CPU0:router(config-bgp)# bgp as-path-loopcheck
```

bgp attribute-download

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

bgp attribute-download

no bgp attribute-download

Syntax Description	This command has no arguments or keywords.

Command Default	BGP attribute download is not enabled.

Command Modes	IPv4 unicast address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines

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

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

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

```
B     100.0.4.0/24 [200/0] via 10.0.101.1, 00:00:37
B     100.0.5.0/24 [200/0] via 10.0.101.1, 00:00:37

RP/0/0/CPU0:router(config)# router bgp 50
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# bgp attribute-download
!
!
!
RP/0/0/CPU0:router# show route bgp

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

bgp auto-policy-soft-reset disable

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

bgp auto-policy-soft-reset disable

no bgp auto-policy-soft-reset disable

Syntax Description   This command has no arguments or keywords.

Command Default   Automatic soft reset of peers is enabled.

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **disable** keyword was changed from optional to mandatory.
Release 3.3.0	This command was supported in VRF configuration mode.

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

Task ID

Task ID	Operations
bgp	read, write

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

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

bgp bestpath as-path ignore

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

bgp bestpath as-path ignore

no bgp bestpath as-path ignore

Syntax Description This command has no arguments or keywords.

Command Default The autonomous system path length is used (not ignored) when a best path is selected.

Command Modes Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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

Task ID

Task ID	Operations
bgp	read, write

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

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

Related Commands

Command	Description
bgp bestpath compare-routerid, page 46	Compares identical routes received from eBGP peers during the best-path selection process and selects the route with the lowest router ID.
bgp bestpath med always, page 50	Allows the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
bgp bestpath med confed, page 52	Enables MED comparison among paths learned from confederation peers.
bgp bestpath med missing-as-worst, page 54	Enables the software to consider a missing MED attribute in a path as having a value of infinity.

bgp bestpath compare-routerid

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

bgp bestpath compare-routerid

no bgp bestpath compare-routerid

Syntax Description   This command has no arguments or keywords.

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

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router(config)# router bgp 500
```

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

Related Commands

Command	Description
show bgp, page 244	Displays entries in the BGP routing table.

bgp bestpath cost-community ignore

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

bgp bestpath cost-community ignore

no bgp bestpath cost-community ignore

Syntax Description This command has no arguments or keywords.

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

Command Modes Router configuration

VRF configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes 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 **bgp bestpath cost-community ignore** command to disable the evaluation of the cost community attribute to help isolate problems and troubleshoot issues that relate to BGP path selection. This command can also be used to delay the activation of cost community attribute evaluation so that cost community filtering can be deployed in a large network at the same time.

Task ID

Task ID	Operations
bgp	read, write

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

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

Related Commands

Command	Description
show bgp, page 244	Displays entries in the BGP routing table.

bgp bestpath med always

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

bgp bestpath med always

no bgp bestpath med always

Syntax Description   This command has no arguments or keywords.

Command Default   The software does not compare MEDs for paths from neighbors in different autonomous systems.

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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 MED is one of the parameters that is considered by the software when selecting the best path among many alternative paths. The software chooses the path with the lowest MED.

By default, during the best-path selection process, the software makes a MED comparison only among paths from the same autonomous system. This command changes the default behavior of the software by allowing comparison of MEDs among paths regardless of the autonomous system from which the paths are received.

When the **bgp bestpath med always** command is not enabled and distributed BGP is configured, speakers calculate partial best paths only (executes the best-path steps up to the MED comparison) and send them to BGP Routing Information Base (bRIB). bRIB calculates the final best path (executes all the steps in the best-path calculation). When the **bgp bestpath med always** command is enabled and distributed BGP is configured, speakers can compare the MED across all ASs, allowing the speaker to calculate a single best path to send it to bRIB. bRIB is the ultimate process that calculates the final best path, but when the **bgp bestpath med always** command is enabled, the speakers send a single best path instead of potentially sending multiple, partial best paths

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
bgp bestpath med confed, page 52	Enables MED comparison among paths learned from confederation peers.
bgp bestpath med missing-as-worst, page 54	Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path.
show bgp, page 244	Displays entries in the BGP routing table.

bgp bestpath med confed

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

bgp bestpath med confed

no bgp bestpath med confed

Syntax Description   This command has no arguments or keywords.

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

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

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

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

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

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

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

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

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

- Path 1: 65000 65004, med = 2, IGP metric = 20

- Path 2: 65001 65004, med = 3, IGP metric = 10

- Path 3: 65002 1, med = 1, IGP metric = 30

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

- Has a lower MED than path 2

- Has a lower IGP metric than path 3

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
bgp bestpath med always, page 50	Enables MED comparison among paths from neighbors in different autonomous systems.
bgp bestpath med missing-as-worst, page 54	Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path.
show bgp, page 244	Displays entries in the BGP routing table.

bgp bestpath med missing-as-worst

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

bgp bestpath med missing-as-worst

no bgp bestpath med missing-as-worst

Syntax Description   This command has no arguments or keywords.

Command Default   The software assigns a value of 0 to the missing MED, causing the path with the missing MED attribute to be considered as the best possible MED.

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to direct the Border Gateway Protocol (BGP) software to consider a missing MED attribute in a path as having a value of infinity, making this path the least desirable path:

```
RP/0/0/CPU0:router(config)# router bgp 210
RP/0/0/CPU0:router(config-bgp)# bgp bestpath med missing-as-worst
```

Related Commands

Command	Description
bgp bestpath med always, page 50	Enables MED comparison among paths from neighbors in different autonomous systems.
bgp bestpath med confed, page 52	Enables MED comparison among paths learned from confederation peers.
show bgp, page 244	Displays entries in the BGP routing table.

bgp client-to-client reflection disable

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

bgp client-to-client reflection [cluster-id *cluster-id]* **disable**

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

Syntax Description

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

Command Default   Client-to-client reflection is enabled.

Command Modes   Address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **disable** keyword was changed from optional to mandatory.
Release 3.8.0	Support was added for multiple cluster-IDs.

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.

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

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

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

```
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit

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

Related Commands

Command	Description
bgp cluster-id, page 58	Configures the cluster ID if the BGP cluster has more than one route reflector.
route-reflector-client, page 228	Configures the router as a BGP route reflector and configures the specified neighbor as its client.
show bgp, page 244	Displays entries in the BGP routing table.

bgp cluster-id

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

bgp cluster-id *cluster-id*

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

Syntax Description

cluster-id	Cluster ID of this router acting as a route reflector; maximum of 4 bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295.

Command Default   A cluster ID is not configured.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

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

Together, a route reflector and its clients form a *cluster*. A cluster of clients usually has a single route reflector. In such instances, the cluster is identified by the software as the router ID of the route reflector. To increase redundancy and avoid a single point of failure in the network, a cluster might have more than one route reflector. If it does, all route reflectors in the cluster must be configured with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster.

A single route reflector can also support multiple clusters. Each cluster is identified by a unique cluster-id. The cluster-id configured by the **bgp cluster-id** command is taken as the default. If bgp cluster-id is not configured, the router ID for the default VRF identifies the default cluster. A neighbor can be associated with one cluster only, and the corresponding cluster-id is configured in neighbor configuration mode. If the cluster-id is not configured for a neighbor and the neighbor is a route reflector client, then the neighbor is assigned to the default cluster.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the local router as one of the route reflectors serving the cluster. Neighbor 192.168.70.24 is assigned to the default cluster with cluster-id 1.

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

Related Commands

Command	Description
cluster-id	Configures the cluster to which a neighbor belongs.
route-reflector-client, page 228	Configures the router as a BGP route reflector and configures the specified neighbor as its client.
show bgp, page 244	Displays entries in the BGP routing table.

bgp confederation identifier

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

bgp confederation identifier *as-number*

no bgp confederation identifier [*as-number*]

Syntax Description

as-number	Autonomous system (AS) number that internally includes multiple autonomous systems. • 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.

Command Default

No confederation identifier is configured.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

One way to reduce the internal BGP (iBGP) mesh is to divide an autonomous system into multiple autonomous systems and group them into a single confederation. Each autonomous system is fully meshed within itself, and has a few connections to another autonomous system in the same confederation. Although the peers in different autonomous systems have external BGP (eBGP) sessions, they exchange routing information as if they are iBGP peers. Specifically, the confederation maintains the next hop and local preference information, and that allows you to retain a single Interior Gateway Protocol (IGP) for all autonomous systems. To the outside world, the confederation looks like a single autonomous system.

Use the **bgp confederation identifier** command to specify the autonomous system number for the confederation. This autonomous system number is used when BGP sessions are established with external peers in autonomous systems that are not part of the confederation.

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
bgp confederation peers, page 62	Configures the autonomous systems that belong to the confederation.

bgp confederation peers

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

bgp confederation peers [*as-number*]

no bgp confederation peers [*as-number*]

Syntax Description

as-number	Autonomous system (AS) numbers for Border Gateway Protocol (BGP) peers that belong to the confederation. • 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.

Command Default

No BGP peers are identified as belonging to the confederation.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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 autonomous systems specified in this command are visible internally to a confederation. Each autonomous system is fully meshed within itself. The bgp confederation identifier, page 60 command specifies the confederation to which the autonomous systems belong.

To specify multiple autonomous systems, enter BGP confederation peer configuration mode then enter one *autonomous-system-number* for each command line.

Task ID	Task ID	Operations
	bgp	read, write

Examples   The following example shows that autonomous systems 1090 and 1093 belong to a single confederation:

```
RP/0/0/CPU0:router(config)# router bgp 1090
RP/0/0/CPU0:router(config-bgp)# bgp confederation peers 1093
```

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

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

Related Commands

Command	Description
bgp confederation identifier, page 60	Specifies a BGP confederation identifier.

bgp dampening

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

bgp dampening [*half-life* [*reuse suppress max-suppress-time*]| **route-policy** *route-policy-name*]

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

Syntax Description

half-life	(Optional) Time (in minutes) after which a penalty is decreased. Once the route has been assigned a penalty, the penalty is decreased by half after the half-life period (which is 15 minutes by default). Penalty reduction happens every 5 seconds. Range of the half-life period is from 1 to 45 minutes.
reuse	(Optional) Value for route reuse if the flapping route penalty decreases and falls below the reuse value. When this happens, the route is unsuppressed. The process of unsuppressing routes occurs at 10-second increments. Range is 1 to 20000.
suppress	(Optional) Maximum penalty value. Suppress a route when its penalty exceeds the value specified. When this happens, the route is suppressed. Range is 1 to 20000.
max-suppress-time	(Optional) Maximum time (in minutes) a route can be suppressed. Range is 1 to 255. If the *half-life* value is allowed to default, the maximum suppress time defaults to 60 minutes.
route-policy *route-policy-name*	(Optional) Specifies the route policy to use to set dampening parameters.

Command Default

Route dampening is disabled.

half-life : 15 minutes

reuse : 750

suppress : 2000

max-suppress-time : four times *half-life* value

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VRF IPv4 address family configuration

VPNv6 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **policy** keyword was changed to **route-policy** .
Release 3.3.0	This command was supported in the following configuration modes:
	• VPNv4 address family
	• VRF IPv4 address family
Release 3.5.0	This command was supported in the following configuration modes:
	• VPNv6 address family
	• VRF IPv6 address family

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
clear bgp dampening, page 109	Clears BGP route dampening information and unsuppresses the suppressed routes.
clear bgp flap-statistics, page 114	Clears BGP flap statistics.
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

Command	Description
show bgp dampened-paths, page 286	Displays BGP dampened routes.
show bgp flap-statistics, page 291	Displays BGP flap statistics.
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.

bgp default local-preference

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

bgp default local-preference *value*

no bgp default local-preference [*value*]

Syntax Description

value	Local preference value. Range is 0 to 4294967295. Higher values are preferable.

Command Default   Enabled with a value of 100.

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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.

Generally, the default value of 100 allows you to easily define a particular path as less preferable than paths with no local preference attribute. The preference is sent to all networking devices in the local autonomous system.

Task ID

Task ID	Operations
bgp	read, write

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

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

bgp enforce-first-as disable

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

bgp enforce-first-as disable

no bgp enforce-first-as disable

Syntax Description   This command has no arguments or keywords.

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

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **disable** keyword was changed from optional to mandatory.
Release 3.3.0	This command was supported in VRF configuration mode.

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.

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

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# bgp enforce-first-as disable
```

Related Commands

Command	Description
enforce-first-as, page 148	Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode.
enforce-first-as-disable, page 150	Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode.
show bgp, page 244	Displays entries in the BGP routing table.

bgp fast-external-fallover disable

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

bgp fast-external-fallover disable

no bgp fast-external-fallover disable

Syntax Description

disable	Disables BGP fast external failover.

Command Default

BGP sessions of any directly adjacent external peers are immediately reset if the link used to reach them goes down.

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **disable** keyword was changed from optional to mandatory.
Release 3.3.0	This command was supported in VRF configuration mode.

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.

By default, BGP sessions of any directly adjacent external peers are immediately reset, which allows the network to recover faster when links go down between BGP peers.

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

bgp graceful-restart

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

bgp graceful-restart

no bgp graceful-restart

Syntax Description   This command has no arguments or keywords.

Command Default   Graceful restart support is not enabled.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 **bgp graceful-restart** command to enable graceful restart functionality on the router, and also to advertise graceful restart to neighboring routers.

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

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples The following example shows how to enable graceful restart:

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

Related Commands

Command	Description
bgp graceful-restart graceful-reset, page 73	Enables a graceful reset if configuration changes force a peer reset.
bgp graceful-restart purge-time, page 75	Defines the maximum time before stale routes are purged.
bgp graceful-restart restart-time, page 77	Defines the maximum time advertised to neighbors
bgp graceful-restart stalepath-time, page 79	Defines the maximum time to wait for the End-of-RIB message from a neighbor that has been restarted before deleting learned routes.
show bgp, page 244	Displays entries in the BGP routing table.
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.
show bgp process, page 357	Displays BGP process information.

bgp graceful-restart graceful-reset

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

bgp graceful-restart graceful-reset

no bgp graceful-restart graceful-reset

Syntax Description   This command has no arguments or keywords.

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

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable graceful reset:

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

Related Commands

Command	Description
bgp graceful-restart, page 71	Enables a graceful restart.
show bgp, page 244	Displays entries in the BGP routing table.

Command	Description
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.
show bgp process, page 357	Displays BGP process information.

bgp graceful-restart purge-time

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

bgp graceful-restart purge-time *seconds*

no bgp graceful-restart purge-time *seconds*

Syntax Description

| *seconds* | Maximum time before stale routes are purged. Time in seconds. Range is 0 to 6000. |

Command Default

seconds : 600

Command Modes

Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
bgp graceful-restart, page 71	Enables a graceful restart.

Command	Description
show bgp, page 244	Displays entries in the BGP routing table.
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.
show bgp process, page 357	Displays BGP process information.

bgp graceful-restart restart-time

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

bgp graceful-restart restart-time *seconds*

no bgp graceful-restart restart-time *seconds*

Syntax Description

seconds	Maximum time advertised to neighbors. Time in seconds. Range is 1 to 4095.

Command Default

seconds : 120

Command Modes

Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to change the BGP graceful restart time to 400 seconds:

```
RP/0/0/CPU0:router(config)#router bgp 3
RP/0/0/CPU0:router(config-bgp)# bgp graceful-restart restart-time 400
```

Related Commands

Command	Description
bgp graceful-restart, page 71	Enables a graceful restart.

Command	Description
show bgp, page 244	Displays entries in the BGP routing table.
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.
show bgp process, page 357	Displays BGP process information.

bgp graceful-restart stalepath-time

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

bgp graceful-restart stalepath-time *seconds*

no bgp graceful-restart stalepath-time *seconds*

Syntax Description

seconds	Maximum wait time. Time in seconds. Range is 1 to 4095.

Command Default   *seconds* : 360

Command Modes   Router configuration

Command History

Release	Modifications
Release 3.2	This command was introduced.

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

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

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

Task ID

Task ID	Operations
bgp	read, write

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

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

Related Commands

Command	Description
bgp graceful-restart, page 71	Enables a graceful restart.
show bgp, page 244	Displays entries in the BGP routing table.
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.
show bgp process, page 357	Displays BGP process information.

bgp import-delay

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

bgp import-delay *seconds milliseconds*

no bgp import-delay

Syntax Description

seconds	Specifies batch import processing delay in seconds. Range is 0 to 10 seconds.
milliseconds	Specifies batch import processing delay in milliseconds. Range is 0 to 999 seconds.

Command Default No delay is configured.

Command Modes Address-family VPNv4 Unicast

Address-family VPNv6 Unicast

Command History

Release	Modification
Release 3.9.1	This command was introduced.

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

Task ID

Task ID	Operation
bgp	read, write

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

```
RP/0/0/CPU0:router#configure
RP/0/0/CPU0:router(config)#router bgp 100
RP/0/0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/0/CPU0:router(config-bgp-af)#bgp import-delay 2 0
```

Related Commands

Command	Description
bgp label-delay, page 83	Enables delay for Border Gateway Protocol (BGP) batch label processing

bgp label-delay

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

bgp label-delay *seconds milliseconds*

no bgp label-delay

Syntax Description		
seconds	Specifies batch label processing delay in seconds. Range is 0 to 10 seconds.
milliseconds	Specifies batch label processing delay in milliseconds. Range is 0 to 999 seconds.

Command Default   No delay is configured.

Command Modes   Address-family IPv4 Unicast

Address-family IPv4 Multicast

Address-family IPv6 Unicast

Address-family IPv6 Multicast

Address-family VPNv4 Unicast

Address-family VPNv6 Unicast

Command History

Release	Modification
Release 3.9.1	This command was introduced.

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

Task ID

Task ID	Operation
bgp	read, write

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

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

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

Related Commands

Command	Description
bgp import-delay, page 81	Enables delay for Border Gateway Protocol (BGP) batch import processing

bgp log neighbor changes disable

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

bgp log neighbor changes disable

no bgp log neighbor changes disable

Syntax Description   This command has no arguments or keywords.

Command Default   BGP neighbor changes are logged.

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **disable** keyword was changed from optional to mandatory.
Release 3.3.0	This command was supported in VRF configuration mode.

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.

Logging of BGP neighbor status changes (up or down) and resets is used for troubleshooting network connectivity problems and measuring network stability. Unexpected neighbor resets might indicate high error rates or high packet loss in the network, and should be investigated.

Status change message logging does not substantially affect performance, unlike, for example, enabling per-BGP update debugging. If the UNIX syslog facility is enabled, messages are sent by the software to the UNIX host running the syslog daemon so that the messages can be stored and archived on disk. If the UNIX syslog facility is not enabled, the status change messages are kept in the internal buffer of the router, and are not stored to disk.

The neighbor status change messages are not tracked if the **bgp log neighbor changes disable** command is disabled, except for the last reset reason, which is always available as output of the **show bgp neighbors** command.

Up and down messages for BGP neighbors are logged by the software by default. Use the **bgp log neighbor changes disable** command to stop logging BGP neighbor changes.

Task ID

Task ID	Operations
bgp	read, write

Examples

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

```
RP/0/0/CPU0:router(config)# router bgp 65530
RP/0/0/CPU0:router(config-bgp)# bgp log neighbor changes disable
```

Related Commands

Command	Description
show bgp neighbors, page 310	Displays information about the TCP and BGP connections to neighbors.

bgp maximum neighbor

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

bgp maximum neighbor *limit*

no maximum neighbor [*limit*]

Syntax Description

limit	Maximum number of neighbors. Range is 1 to 15000.

Command Default   Default limit is 4000

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

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

Any attempt to configure the neighbor limit below 1 or above 15000 fails. Similarly, attempting to configure the limit below the number of neighbors currently configured fails. For example, if there are 3250 neighbors configured, you cannot set the *limit* below 3250.

Task ID

Task ID	Operations
bgp	write

Examples   The following example shows how to change the default maximum neighbor limit and set it to 1200:

```
RP/0/0/CPU0:router(config)#router bgp 65530
RP/0/0/CPU0:router(config-bgp)# bgp maximum neighbor 1200
```

bgp redistribute-internal

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

bgp redistribute-internal

no bgp redistribute-internal

Syntax Description This command has no arguments or keywords.

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

Command Modes Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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

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

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router(config)#router bgp 1
```

```
RP/0/0/CPU0:router(config-bgp)# bgp redistribute-internal
RP/0/0/CPU0:router(config-bgp)# exit
RP/0/0/CPU0:router(config)# router ospf area1
RP/0/0/CPU0:router(config-router)# redistribute bgp 1
RP/0/0/CPU0:router(config-router)# end
RP/0/0/CPU0:router# clear route *
```

Related Commands

Command	Description
clear bgp, page 104 *	Resets all BGP neighbors.
clear route *	Resets all routes.

bgp router-id

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

bgp router-id *ip-address*

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

Syntax Description

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

Command Default

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

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode. The *interface-type interface-instance* arguments were removed.

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.

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

For more details on router IDs, see the *Cisco IOS XR Routing Configuration Guide for the Cisco XR 12000 Series Router*

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

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

Related Commands

Command	Description
show bgp, page 244	Displays entries in the BGP routing table.

bgp scan-time

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

bgp scan-time *seconds*

no bgp scan-time *seconds*

Syntax Description

seconds	Scanning interval (in seconds) of BGP routing information. Range is 5 to 3600 seconds.

Command Default The default scanning interval is 60 seconds.

Command Modes Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VPNv4 address family configuration mode.
Release 3.5.0	This command was supported in VPNv6 address family configuration mode.
Release 4.0.0	Support was removed for all address family configuration modes.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

This example shows how to set the scanning interval to 20 seconds:

```
RP/0/0/CPU0:router(config)# router bgp 64500
RP/0/0/CPU0:router(config-bgp-af)# bgp scan-time 20
```

Related Commands

Command	Description
show bgp, page 244	Displays entries in the BGP routing table.

bgp update-delay

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

bgp update-delay *seconds* **[always]**

nobgp update-delay [*seconds*][**always**]

Syntax Description

seconds	Delay in seconds for the router to send the first updates. Range is 0 to 3600.
always	(Optional) Specifies that the router always wait for the update delay time, even if all neighbors have finished sending their initial updates sooner.

Command Default

120 seconds

Command Modes

Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines

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

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

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

Task ID

Task ID	Operations
bgp	read, write

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

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

bgp write-limit

To modify the upper bounds on update message queue lengths or to enable desynchronization, use the **bgp write-limit** command in an appropriate configuration mode. To return the bounds to their default values and to disable desynchronization, use the **no** form of this command.

bgp write-limit *group-limit global-limit* **[desynchronize]**

no bgp write-limit [*group-limit global-limit*] **[desynchronize]**

Syntax Description

group-limit	Per-update group limit on the number of update messages the software queues. Range is 500 to 100000000. Group limit cannot be greater than the global limit.
global-limit	Global limit on the number of update messages the software queues. Range is 500 to 100000000.
desynchronize	(Optional) Enables desynchronization.

Command Default

group-limit : 50,000

global-limit : 250,000

Desynchronizationis off.

Command Modes

Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The *group-limit* and *global-limit* default values have changed.

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 **bgp write-limit** command to configure both a per-update group and a global limit on the number of messages the software queues when updating peers. Increasing these limits can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory use during convergence. In addition, this command can be used to enable desynchronization. Desynchronization can decrease memory use and speed up convergence for the fastest neighbors if one or more neighbors in an update group process updates significantly slower than other neighbors in the same group. However, enabling desynchronization can cause a significant degradation in overall convergence time, especially if the router is experiencing high CPU utilization. For this reason, enabling desynchronization is discouraged.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure BGP to operate with a per-update group limit of 9000 messages and a global limit of 27,000 messages:

```
RP/0/0/CPU0:router(config)# router bgp 65000
RP/0/0/CPU0:router(config-bgp)#bgp write-limit 9000 27000
```

capability orf prefix

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

capability orf prefix {**receive**| **send**| **both**| **none**}

no capability orf prefix [**receive**| **send**| **both**| **none**]

Syntax Description

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

Command Default

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

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

VRF neighbor IPv4 address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was changed from **capability orf prefix-list** to **capability orf prefix** . This command was supported in VRF neighbor IPv4 address family configuration mode.

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 advertisement of the prefix list ORF capability by a BGP speaker indicates whether the speaker can send prefix lists to the specified neighbor and whether it accepts prefix lists from the neighbor. The speaker sends

a prefix list if it indicated the ability to send them, and if the neighbor indicated it was willing to accept them. Similarly, the neighbor sends a prefix list to the speaker if it indicated the ability to send them and the speaker indicated the willingness to accept them.

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

If the neighbor sends a prefix list and the speaker accepts it, the speaker applies the received prefix list, plus any locally configured outbound filters, to limit its outbound routing updates to the neighbor. Increased filtering prevents unwanted routing updates between neighbors and reduces resource requirements for routing update generation and processing.

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

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
show bgp neighbors, page 310	Displays information about BGP neighbors. Use the **received prefix-filter** keywords to display information on the prefix list filter.

capability suppress 4-byte-as

To suppress 4-byte AS capability from being advertised to the BGP peer, use the **capability suppress 4-byte-as** command in the appropriate configuration mode. To remove the **capability suppress 4-byte-as** command from the configuration and restore the system to the default condition, in which the software advertises the capability, use the **no** form of this command.

capability suppress 4-byte-as [**inheritance-disable**]

no capability suppress 4-byte-as

Syntax Description

inheritance-disable	Prevents capability suppress 4-type-as being inherited from the parent.

Command Default

4-byte-as capability is advertised to the BGP peer.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.4.1	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

By default, the software advertises the 4-byte AS capability to BGP peers. To override this default behavior, use the **capability suppress 4-byte-as** command under the command modes listed in the "Command Modes" section. If configured under the neighbor group or session group, all neighbors using the group inherit the configuration. Use the **no** option to remove the command.

⚠

Caution   The BGP session resets automatically, if the 4-byte AS capability of an existing BGP session is changed by configuring **capability suppress 4-byte-as** or **capability suppress 4-byte-as** **inheritance-disable** .

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to configure the **capability suppress 4-byte-as** command:

```
RP/0/0/CPU0:router# show bgp nei 10.3.3.3 conf
neighbor 10.3.3.3
 remote-as 65000             [n:internal]
 description PE3             []
 update-source Loopback0     [n:internal]
 address-family ipv4 unicast [n:internal]

RP/0/0/CPU0:router#show bgp nei 10.3.3.3
BGP neighbor is 10.3.3.3
 Remote AS 65000, local AS 65000, internal link
 Description: PE3
 Remote router ID 10.3.3.3
  BGP state = Established, up for 1w0d
  Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 25962 messages, 0 notifications, 0 in queue
  Sent 25968 messages, 1 notifications, 0 in queue
  Minimum time between advertisement runs is 0 seconds

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

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

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

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

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

```
     Minimum time between advertisement runs is 0 seconds

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

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

With the **inheritance-disable** keyword:

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

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

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

clear bgp

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

clear bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast**| **all**| **tunnel**| **mdt**}| **ipv6** {**unicast**| **multicast**| **all**| **labeled-unicast**}| **all** {**unicast**| **multicast**| **all**| **labeled-unicast**| **mdt**| **tunnel**}| **vpnv4 unicast**| **vrf** {*vrf-name*| **all**} {**ipv4** {**unicast**| **labeled-unicast**}| **ipv6 unicast**}| **vpnv6 unicast**]

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.
mdt	(Optional) Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 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).
vrf-name	Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) For VRF, specifies IPv4 unicast and labeled-unicast address families.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address prefixes.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address prefixes.

Command Default   No default behavior or values

Command Modes   EXEC

Command History	**Release**	**Modification**
	Release 3.2	This command was introduced.
	Release 3.3.0	The following keywords and argument were added: • **vpnv4 unicast** • **vrf** • *vrf-name* • **all** • **ipv4** { **unicast** | **labeled-unicast** }
	Release 3.4.0	The **as** keyword has been added and the *as-number* 4-byte number range 1.0 to 65535.65535 was supported. The following keywords were added: • **ipv4 multicast** • **ipv4 all** • **ipv6 all** • **ipv6 unicast** • **ipv6 multicast** • **soft** • **labeled-unicast**
	Release 3.5.0	The following keywords were added: • **tunnel** • **mdt** • **ipv6 unicast** • **vpnv6 unicast** The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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

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

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to hard reset neighbor 10.0.0.1:

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

Related Commands

Command	Description
clear bgp self-originated, page 124	Clears self-originated routes.
clear bgp soft, page 128	Soft resets a group of BGP neighbors.
show bgp, page 244	Displays entries in the BGP routing table.
show bgp neighbors, page 310	Displays information about the TCP and BGP connections to neighbors.

clear bgp current-mode

To switch from one BGP mode to another, use the **clear bgp current-mode** command in EXEC mode.

clear bgp current-mode

Syntax Description   This command has no arguments or keywords.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.3.0	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 **clear bgp current-mode** command to switch from standalone to distributed mode, or from distributed to standalone mode. The **show bgp process** command indicates the current BGP mode.

> **Note**   Switching from one mode to another causes all BGP sessions to go down.

Task ID

Task ID	Operations
bgp	execute

Examples   The following example shows the **show bgp process** command output before and after switching from one BGP mode to another:

```
RP/0/0/CPU0:router# show bgp process

  BGP Process Information
  BGP is operating in STANDALONE mode
  Autonomous System: 3
  Router ID: 10.18.18.11
  Cluster ID: 10.18.18.11
  Fast external fallover enabled
  Neighbor logging is enabled
  Enforce first AS enabled
  Default local preference: 100
```

```
                    Default keepalive: 60
                    Update delay: 120
                    Generic scan interval: 60

                    Address family: IPv4 Unicast
                    Dampening is not enabled
                    Client reflection is enabled
                    Scan interval: 60
                    Main Table Version: 3
                    IGP notification: IGPs notified
                    RIB has converged: version 0

                    Node             Process    Nbrs Estb Rst Upd-Rcvd Upd-Sent Nfn-Rcv Nfn-Snt
                    node0_0_CPU0     Speaker       5    5  51        0        7       0       5

                    RP/0/0/CPU0:router# configure
                    RP/0/0/CPU0:router(config)# router bgp 3
                    RP/0/0/CPU0:router(config-bgp)# distributed speaker 1
                    RP/0/0/CPU0:router(config-bgp)# distributed speaker 2
                    RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.101.1
                    RP/0/0/CPU0:router(config-bgp-nbr)# speaker-id 2
                    RP/0/0/CPU0:router(config-bgp-nbr)# commit
                    RP/0/0/CPU0:router# clear bgp current-mode
                    RP/0/0/CPU0:router# show bgp process

                    BGP Process Information
                    BGP is operating in DISTRIBUTED mode
                    Autonomous System: 3
                    Router ID: 10.18.18.11
                    Cluster ID: 10.18.18.11
                    Fast external fallover enabled
                    Neighbor logging is enabled
                    Enforce first AS enabled
                    Default local preference: 100
                    Default keepalive: 60
                    Update delay: 120
                    Generic scan interval: 60

                    Address family: IPv4 Unicast
                    Dampening is not enabled
                    Client reflection is enabled
                    Scan interval: 60
                    Main Table Version: 1
                    IGP notification: IGPs not notified
                    RIB has not converged: version 0

                    Node             Process    Nbrs Estb Rst Upd-Rcvd Upd-Sent Nfn-Rcv Nfn-Snt
                    node0_0_CPU0     Speaker 1     4    1  52        0        0       0       4
                    node0_0_CPU0     Speaker 2     1    0   9        0        0       0       1
                    node0_0_CPU0     bRIB 1        0    0   0        0        0       0       0
                    node0_0_CPU0     bRIB 2        0    0   0        0        0       0       0
```

Related Commands

Command	Description
show bgp process, page 357	Displays the current BGP process information.

clear bgp dampening

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

clear bgp dampening

Syntax Description	
ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast and labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
ip-address	(Optional) IP address of the network about which to clear dampening information.
/mask-length	(Optional) Network mask applied to the IP address.

Command Default   If no IP address is specified, dampening information for all routes is cleared.

Command Modes   EXEC

Command History	Release	Modification
	Release 3.2	This command was introduced.
	Release 3.3.0	The following keywords and argument were added: • **vpnv4 unicast** • **vrf** • *vrf-name* • **all** • **ipv4** { **unicast** | **labeled-unicast** }
	Release 3.4.0	The **labeled-unicast** keyword was supported.
	Release 3.5.0	The following keywords were added: • **ipv6 unicast** • **vpnv6 unicast** The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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.

Always use the **clear bgp dampening** command for an individual address-family. The **all** option for address-families with clear bgp dampening should never be used during normal functioning of the system. For example, use

```
clear bgp ipv4 unicast dampening prefix x.x.x./y
```

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear the route dampening information for all 172.20.0.0/16 IPv4 multicast paths:

```
RP/0/0/CPU0:router# clear bgp ipv4 multicast dampening 172.20.0.0/16
```

Related Commands

Command	Description
bgp dampening, page 64	Enables BGP route dampening or changes various BGP route dampening factors.
show bgp dampened-paths, page 286	Displays BGP dampened routes.

clear bgp external

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

clear bgp external

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.
ipv6	(Optional) Specifies IP Version 6 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).
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
graceful	(Optional) Clears all external peers with a hard reset and a graceful restart. This option is available when an address family is not specified.

Command Default   No default behavior or value

Command Modes   EXEC

Command History

Release	Modification
Release 3.5.0	This command was introduced.

Usage Guidelines

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

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear all BGP external peers:

```
RP/0/0/CPU0:router# clear bgp external
```

clear bgp flap-statistics

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

clear bgp flap-statistics

Syntax Description		
ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
regexp *regexp*	(Optional) Clears flap statistics for routes whose AS paths match the regular expression.
route-policy *route-policy-name*	(Optional) Clears flap statistics for the specific route policy.
network	(Optional) Network for which flap counts are to be cleared.
/mask-length	(Optional) Network mask of the network for which flap counts are to be cleared.
ip-address	(Optional) Neighbor address. Clears only flap statistics for routes received from this neighbor.

Command Default   No default behavior or value

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **filter-list** *access-list* keyword and argument were changed to **route-policy** *route-policy-name.*
Release 3.3.0	The following keywords and argument were added:
	• **vpnv4 unicast**
	• **vrf**
	• *vrf-name*
	• **all**
	• **ipv4** { **unicast** | **labeled-unicast** }
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The following keywords were added:
	• **ipv6 unicast**
	• **vpnv6 unicast**
	The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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.

Task ID

Task ID	Operations
bgp	execute

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

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

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

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

clear bgp nexthop performance-statistics

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

clear bgp {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} vpnv6 unicast} nexthop performance-statistics

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
tunnel	Specifies tunnel address prefixes.
mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast**	**labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.4.0	This command was introduced.
Release 3.5.0	The following keywords were added:
	• **tunnel**
	• **mdt**
	• **ipv6 unicast**
	• **vpnv6 unicast**
	The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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 **clear bgp nexthop performance-statistics** command to reset the total number of notifications received from the Routing Information Base (RIB) and the cumulative next-hop processing time. The following information is cleared from the **show bgp nexthops** command output:

- Total critical notifications received
- Total noncritical notifications received
- Best path deleted after last walk
- Best path changed after last walk
- Next-hop table total number of critical and noncritical notifications (Notf) and the time of the last notification received from the RIB (LastRIB) columns (only entries that have a status of unreachable [UR])

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear next-hop performance statistics:

```
RP/0/0/CPU0:router# clear bgp vrf vrf_A nexthop performance statistics
```

Related Commands

Command	Description
show bgp nexthops, page 332	Displays information about the BGP next-hop notifications.

clear bgp nexthop registration

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

clear bgp {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast | vpnv6 unicast} nexthop registration nexthop-address *nexthop-address*

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled-unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
tunnel	Specifies tunnel address prefixes.
mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast**	**labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
nexthop-address	Address of the next hop.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.4.0	This command was introduced.
Release 3.5.0	The following keywords were added: • **tunnel** • **mdt** • **ipv6 unicast** • **vpnv6 unicast** The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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 **clear bgp nexthop registration** command to perform an asynchronous registration of the next hop with the RIB. The **show bgp nexthops** command output shows a critical notification as the LastRIBEvent for the next hop when the **clear bgp nexthop registration** command is used.

Task ID

Task ID	Operations
bgp	execute

Examples   The following example shows how to reregister the next hop with the RIB:

```
RP/0/0/CPU0:router# clear bgp nexthop registration 10.1.1.1
```

Related Commands

Command	Description
show bgp nexthops, page 332	Displays information about the BGP next-hop notifications.

clear bgp peer-drops

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

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

Syntax Description

*	Specifies all BGP neighbors.
ip-address	IP address of a specific network neighbor.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines

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

Task ID

Task ID	Operations
bgp	execute

Examples

The following example shows how to clear the connection-dropped counter for all BGP neighbors:

```
RP/0/0/CPU0:router# clear bgp peer-drops *
```

Related Commands

Command	Description
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.

clear bgp performance-statistics

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

clear bgp [**vrf** {*vrf-name*| **all**}] **performance-statistics**

Syntax Description

vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.

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

Task ID

Task ID	Operations
bgp	execute

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

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

clear bgp self-originated

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

clear bgp {**ipv4**{**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all** } | **all** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **vpnv4 unicast** | **vrf** {*vrf-name* | **all**} | **vpnv6 unicast**} **self-originated**

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast**	**labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.

Command Default No default behavior or values

Command Modes EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.

Release	Modification
Release 3.3.0	The following keywords and argument were added: • **vrf** • *vrf-name* • **all** • **ipv4** { **unicast** | **labeled-unicast** }
Release 3.4.0	The **vpnv4 unicast** keywords were added. The **labeled-unicast** keyword was supported.
Release 3.5.0	The following keywords were added: • **ipv6 unicast** • **vpnv6 unicast** The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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.

Self-originated routes are routes locally originated by the **network** command, **redistribute** command, or **aggregate-address** command.

Task ID

Task ID	Operations
bgp	execute

Examples   The following example shows how to clear self-originated IPv4 routes:

```
RP/0/0/CPU0:router# clear bgp ipv4 unicast self-originated
```

Related Commands

Command	Description
aggregate-address, page 24	Creates an aggregate entry in a BGP routing table.
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
redistribute (BGP), page 211	Redistributes routes from another routing protocol into BGP.

clear bgp shutdown

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

clear bgp {**ipv4**{**unicast** | **multicast** | **labeled-unicast** | **all**} | **ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all** } | **all** {**unicast** | **multicast** | **labeled-unicast** | **all**} | **vpnv4 unicast** | **vrf** {*vrf-name* | **all**} | **vpnv6 unicast**} **shutdown**

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast**	**labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.

Release	Modification
Release 3.3.0	The following keywords and argument were added: • **vpnv4 unicast** • **vrf** • *vrf-name* • **all** • **ipv4** { **unicast** | **labeled-unicast** }
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The following keywords were added: • **ipv6 unicast** • **vpnv6 unicast** The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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.

Task ID

Task ID	Operations
bgp	execute

Examples

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

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

Related Commands

Command	Description
show bgp, page 244	Displays entries in the BGP routing table.
show bgp neighbors, page 310	Displays information about the TCP and BGP connections to neighbors.

clear bgp soft

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

clear bgp {**ipv4**{**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **ipv6** {**unicast** | **multicast** | **labeled-unicast** | **all** } | **all** {**unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | **mdt**} | **vpnv4 unicast** | **vrf** {*vrf-name* | **all**} | **vpnv6 unicast**} {***** | *ip-address* | **as**| *as-number* | **external**}**soft**[[**in**| {**prefix-filter**} | **out**]]

Syntax Description

ipv4	Specifies IP Version 4 address prefixes.
unicast	Specifies unicast address prefixes.
multicast	Specifies multicast address prefixes.
labeled-unicast	Specifies labeled unicast address prefixes.
all	For subaddress families, specifies prefixes for all subaddress families.
tunnel	Specifies tunnel address prefixes.
mdt	Specifies IPv4 multicast distribution tree (MDT) address prefixes.
ipv6	Specifies IP Version 6 address prefixes.
all	For address family, specifies prefixes for all address families.
vpnv4 unicast	Specifies VPNv4 unicast address families.
vrf	Specifies VPN routing and forwarding (VRF).
vrf-name	Name of a VRF.
all	For VRF, specifies all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	For VRF, specifies IPv4 unicast or labeled-unicast address families.
ipv6 unicast	For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	Specifies VPNv6 unicast address families.
*	Soft resets all BGP neighbors.
ip-address	IP address of the neighbor to be reset.
as *as-number*	Autonomous system (AS) number for all neighbors to be reset. Range for 2-byte numbers is 1 to 65535. Range for 4-byte numbers is 1.0 to 65535.65535.

external	Specifies clearing of all external peers.
in	(Optional) Triggers an inbound soft reset. If the **in** or **out** keyword is not specified, both inbound and outbound soft resets are triggered.
prefix-filter	(Optional) Specifies to send a new Outbound Route Filter (ORF) to the neighbor. Neighbor installs the new ORF and resends its routes.
out	(Optional) Triggers an outbound soft reset. If the **in** or **out** keyword is not specified, both inbound and outbound soft resets are triggered.

Command Default   No default behavior or value

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The following keywords and argument were added: • **vpnv4 unicast** • **vrf** • *vrf-name* • **all** • **ipv4** { **unicast** | **labeled-unicast** }
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported. The **labeled-unicast** keyword was supported.
Release 3.5.0	The following keywords were added: • **tunnel** • **ipv6 unicast** • **vpnv6 unicast** The **labeled-unicast** keyword was supported for **ipv6** and **all** address families.

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

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

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

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

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

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

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

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

Task ID

Task ID	Operations
bgp	execute

Examples

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

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

Related Commands

Command	Description
bgp auto-policy-soft-reset disable, page 42	Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified.
clear bgp, page 104	Resets a group of BGP neighbors.
clear bgp self-originated, page 124	Clears self-originated routes.
show bgp, page 244	Displays entries in the BGP routing table.

Command	Description
show bgp neighbors, page 310	Displays information about the TCP and BGP connections to neighbors.
soft-reconfiguration inbound, page 439	Configures the software to store updates received from a neighbor.

default-information originate (BGP)

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

default-information originate

no default-information originate

Syntax Description   This command has no arguments or keywords.

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

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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

Task ID

Task ID	Operations
bgp	read, write

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

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

Related Commands

Command	Description
redistribute (BGP), page 211	Redistributes routes from another protocol into BGP.

default-metric (BGP)

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

default-metric *value*

no default-metric [*value*]

Syntax Description

value	Default metric value appropriate for the specified routing protocol. Range is 1 to 4294967295.

Command Default   A metric is not set.

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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 **default-metric** command to set the Multi Exit Discriminator (MED) to advertise to peers for routes that do not already have a metric set (routes that were received with no MED attribute).

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to set the BGP default metric:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# default-metric 10
```

default-originate

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

default-originate [**inheritance-disable**| **route-policy** *route-policy-name*]

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

Syntax Description

inheritance-disable	(Optional) Prevents the **default-originate** command characteristics from being inherited from a parent group.
route-policy *route-policy-name*	(Optional) Specifies the name of a route policy. The route policy allows route 0.0.0.0 to be injected conditionally. IPv6 address family is supported.

Command Default

The default route is not advertised to BGP neighbors.

Command Modes

IPv4 neighbor address family configuration

IPv6 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

IPv4 address family group configuration

IPv6 address family group configuration

VRF IPv4 neighbor address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **policy** keyword was changed to **route-policy** .
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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 **default-originate** command does not require the presence of the default route (0.0.0.0/0 for IPv4 or ::/0 for IPv6) in the local router. When the **default-originate** command is used with a route policy, the default route is advertised if any route in the BGP table matches the policy.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to unconditionally advertise the route 0.0.0.0/0 to the neighbor 172.20.2.3:

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

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

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

Related Commands

Command	Description
default-information originate (BGP), page 132	Allows the default route to be redistributed into BGP from another routing protocol.
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.

description (BGP)

To annotate a neighbor, neighbor group, VPN routing and forwarding (VRF) neighbor, or session group, use the **description** command in an appropriate configuration mode. To remove the annotation, use the **no** form of this command.

description *text*

no description [*text*]

Syntax Description

text	Meaningful description or comment. Maximum of 80 characters.

Command Default   No comment or description exists.

Command Modes   Neighbor group configuration

Neighbor configuration

Session group configuration

VRF neighbor configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)#router bgp 65000
```

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

distance bgp

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

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

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

Syntax Description

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

Command Default

external-distance : 20

internal-distance : 200

local-distance : 200

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF IPv4 address family configuration mode.
Release 3.5.0	This command was supported in VRF IPv6 address family configuration mode.

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

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

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

Task ID

Task ID	Operations
bgp	read, write

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

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

Related Commands

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

distributed speaker

To enable a distributed speaker process, use the **distributed speaker** command in router configuration mode. To remove the distributed speaker process, use the **no** form of this command.

distributed speaker *id*

no distributed speaker *id*

Syntax Description

id	ID of the distributed speaker process. Range is 1 to 15.

Command Default   Default is 0.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.

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

If BGP is currently running in standalone mode, you must enter the **clear bgp current-mode** command to switch from standalone or distributed mode.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable distributed speaker process 3:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# distributed speaker 3
```

Related Commands

Command	Description
clear bgp current-mode, page 107	Switch BGP from one mode to another

dmz-link-bandwidth

To originate a demilitarized zone (DMZ) link bandwidth extended community for the link to an eBGP or iBGP neighbor, use the **dmz-link-bandwidth** command in an Neighbor configuration mode. To stop origination of the DMZ link bandwidth extended community, use the **no** form of this command.

dmz-link-bandwidth [**inheritance-disable**]

no dmz-link-bandwidth

Syntax Description

inheritance-disable	(Optional) Prevents the **dmz-link-bandwidth** command from being inherited from a parent group.

Command Default   BGP does not originate the DMZ link bandwidth extended community.

Command Modes   Neighbor configuration

Command History

Release	Modification
Release 4.1.1	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 **dmz-link-bandwidth** command to advertise the bandwidth of links that are used to exit an autonomous system.

Task ID

Task ID	Operations
bgp	read, write

Examples   This example shows how to advertise the bandwidth of links to eBGP neighbors from router bgp 1:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)#neighbor 45.67.89.01
RP/0/0/CPU0:router(config-bgp-nbr)#dmz-link-bandwidth
```

Related Commands

Command	Description
bandwidth	Configures the bandwidth of an interface.
maximum-paths (BGP), page 172	Controls the maximum number of parallel routes that Border Gateway Protocol (BGP) installs in the routing table.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.

dscp (BGP)

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

dscp *value*

no dscp [*value*]

Syntax Description

value	Value of the DSCP. The DSCP value can be a number from 0 to 63, or it can be one of the following keywords: **default** , **ef** , **af11** , **af12** , **af13** , **af21** , **af22** , **af23** , **af31** , **af32** , **af33** , **af41** , **af42** , **af43** , **cs1** , **cs2** , **cs3** , **cs4** , **cs5** , **cs6** , or **cs7** .

Command Default

No default behavior or values

Command Modes

Neighbor configuration

Neighbor session group configuration

Neighbor group configuration

Command History

Release	Modification
Release 3.4.0	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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 **dscp** command to change the minimum and maximum packet thresholds for the DSCP value.

Table 2: dscp Default Settings, page 144 lists the DSCP default settings used by the **dscp** command. The DSCP value, corresponding minimum threshold, maximum threshold, and mark probability are listed. The last row of the table (the row labeled "default") shows the default settings used for any DSCP value not specifically shown in the table.

Table 2: dscp Default Settings

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

DSCP (Precedence)	Minimum Threshold	Maximum Threshold	Mark Probability
af21	32	40	1/10
af22	28	40	1/10
af23	24	40	1/10
af31	32	40	1/10
af32	28	40	1/10
af33	24	40	1/10
af41	32	40	1/10
af42	28	40	1/10
af43	24	40	1/10
cs1	22	40	1/10
cs1	24	40	1/10
cs3	26	40	1/10
cs4	28	40	1/10
cs5	30	40	1/10
cs6	32	40	1/10
cs7	34	40	1/10
ef	36	40	1/10
default	20	40	1/10

Task ID

Task ID	Operations
bgp	read, write

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

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

ebgp-multihop

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

ebgp-multihop [*ttl-value*]

no ebgp-multihop [*ttl-value*]

Syntax Description

ttl-value	(Optional) Time-to-live (TTL) value. Range is 1 to 255 hops.

Command Default   Default TTL value is 255.

Command Modes   Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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 **ebgp-multihop** command to enable multihop peerings with external BGP neighbors. The BGP protocol states that external neighbors must be directly connected (one hop away). The software enforces this by default; however, the **ebgp-multihop** command can be used to override this behavior.

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 allow a BGP connection to neighbor 172.20.16.6 of up to 255 hops away:

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

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.

enforce-first-as

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

enforce-first-as

no enforce-first-as

Syntax Description   This command has no arguments or keywords.

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

Command Modes   Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

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

At any given time, either the **enforce-first-as** command or the enforce-first-as-disable, page 150 command can be configured under a given neighbor, neighbor group or session group. Configuring one command overwrites the other command.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows a configuration in which incoming updates from eBGP neighbors are checked to ensure the first AS number in the AS path is the same as the configured AS number for the neighbor:

```
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/0/CPU0:router(config-bgp-nbr)# enforce-first-as
```

Related Commands

Command	Description
bgp enforce-first-as disable, page 68	Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in router configuration mode and VRF configuration mode.
enforce-first-as-disable, page 150	Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in neighbor configuration mode, neighbor group configuration mode, and session group configuration mode.

enforce-first-as-disable

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

enforce-first-as-disable

no enforce-first-as-disable

Syntax Description   This command has no arguments or keywords.

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

Command Modes   Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

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

At any given time, either the **enforce-first-as-disable** command or the command can be configured under a given neighbor, neighbor group or session group. Configuring one command overwrites the other command.

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/0/CPU0:router(config-bgp-nbr)# enforce-first-as-disable
```

Related Commands

Command	Description
bgp enforce-first-as disable, page 68	Disables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, in router configuration mode and VRF configuration mode.
enforce-first-as, page 148	Enables the software to enforce the first autonomous system in the AS path of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, under neighbor configuration mode, neighbor group configuration mode, and session group configuration mode.

export route-policy

To configure an export route policy, use the **export route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

export route-policy *policy-name*

no export route-policy [*policy-name*]

Syntax Description

policy-name	Name of the configured route policy.

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.5.0	This command was supported in global VRF IPv6 address family configuration mode.

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 **export route-policy** command to define the conditions that allow specified routes to be tagged with specified route-targets.

Task ID

Task ID	Operations
bgp	read, write
ip-services	read, write

Examples

The following example shows how to configure an export route policy:

```
RP/0/0/CPU0:router(config)# vrf vrf-1
RP/0/0/CPU0:router(config-vrf)#address-family ipv4 unicast
RP/0/0/CPU0:router(config-vrf-af)# export route-policy policy-A
```

Related Commands

Command	Description
import route-policy, page 157	Specifies a route policy to import routes into the VRF instance.

export route-target

To configure a VPN routing and forwarding (VRF) export route-target extended community, use the **export route-target** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

export route-target [*as-number:nn* | *ip-address:nn*]

no export route-target [*as-number:nn* | *ip-address:nn*]

Syntax Description

as-number:nn	(Optional) *as-number* —Autonomous system (AS) number of the route-target extended community. • as-number ◦ Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. ◦ Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. • *nn* —32-bit number
ip-address:nn	(Optional) IP address of the route-target extended community. • *ip-address* —32-bit IP address • *nn* —16-bit number

Command Default   No default behavior or values

Command Modes   Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

Export route-target extended communities are associated with prefixes when advertised to remote provider edge (PE) routers. The remote PE routers import the route-target extended communities into a VRF instance that has the import route-targets that match the exported route-target extended communities.

To specify multiple route targets, enter export route target configuration mode then enter one route target for each command line.

Task ID

Task ID	Operations
bgp	read, write
ip-services	read, write

Examples   The following example shows how to specify an export route-target:

```
RP/0/0/CPU0:router(config)# vrf vrf-1
RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-vrf-af)# export route-target 500:1
```

Related Commands

Command	Description
import route-target, page 159	Specifies the import route-target.

ibgp policy out enforce-modifications

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

ibgp policy out enforce-modifications

no ibgp policy out enforce-modifications

Syntax Description   This command has no arguments or keywords.

Command Default   ibgp policy out enforce-modifications is disabled.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 **ibgp policy out enforce-modifications** command to set and modify BGP route attributes for updates to iBGP peers.

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

Task ID

Task ID	Operations
bgp	read, write

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

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

import route-policy

To configure an import route policy, use the **import route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

import route-policy *policy-name*

no import route-policy [*policy-name*]

Syntax Description

policy-name	Name of the configured route policy.

Command Default No default behavior or values

Command Modes Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.5.0	This command was supported in global VRF IPv6 address family configuration mode.

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 **import route-policy** command to define the conditions that allow specified routes to be imported into the VPN routing and forwarding (VRF) instance if the routes are tagged with specified route-targets.

Task ID

Task ID	Operations
bgp	read, write
ip-services	read, write

Examples The following example shows how to allow only policy-B to be imported to VRF:

```
RP/0/0/CPU0:router(config)# vrf vrf-1
RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-vrf-af)# import route-policy policy-B
```

Related Commands

Command	Description
export route-policy, page 152	Specifies a route policy to export routes from the VRF instance.

import route-target

To configure a VPN routing and forwarding (VRF) import route-target extended community, use the **import route-target** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

import route-target [*as-number:nn* | *ip-address:nn*]

noimport route-target [*as-number:nn* | *ip-address:nn*]

Syntax Description

as-number:nn	(Optional) Autonomous system (AS) number of the route-target extended community.
	• *as-number*
	◦ Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
	◦ Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
	◦ Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
	• *nn* —32-bit number
ip-address:nn	(Optional) IP address of the route-target extended community.
	• *ip-address* —32-bit IP address
	• *nn* —16-bit number

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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 **import route-target** command to specify that prefixes associated with the configured import route-target extended communities are imported into the VRF instance.

To specify multiple route targets, enter import route target configuration mode, then enter one route target for each command line.

Task ID

Task ID	Operations
bgp	read, write
ip-services	read, write

Examples

The following example shows how to specify an import route-target:

```
RP/0/0/CPU0:router(config)#vrf vrf-1
RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-vrf-af)# import route-target 500:99
```

Related Commands

Command	Description
export route-target, page 154	Specifies the export route-target.

ignore-connected-check

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

ignore-connected-check [**inheritance-disable**]

no ignore-connected-check

Syntax Description

inheritance-disable	Prevents the **ignore-connected-check** command from being inherited from the parent.

Command Default   Ability to bypass the directly connected next hop check is disabled.

Command Modes   Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable ignore-connected check configuration for neighbor 10.2.3.4:

```
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/0/CPU0:router(config-bgp-nbr)# ignore-connected-check
```

keychain

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

keychain *name*

no keychain [*name*]

Syntax Description

name	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 3.4.0	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# keychain keychain_A
```

Related Commands

Command	Description
keychain-disable, page 164	Overrides any inherited keychain configuration from a neighbor group or session group for BGP neighbors.

keychain-disable

✎

Note Effective with Release 3.9.0, the **keychain-disable** command was replaced by the **keychain inheritance-disable** command. See the keychain inheritance-disable, page 166 command for more information.

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

keychain-disable

no keychain-disable

Syntax Description This command has no arguments or keywords.

Command Default Configured key chains for neighbor and session groups are inherited.

Command Modes Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.4.0	This command was introduced.
Release 3.9.0	This command was replaced by the **keychain inheritance-disable** command.

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.

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

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# session-group group1
RP/0/0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/0/CPU0:router(config-bgp-nbr)#use session-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# keychain-disable
```

Related Commands

Command	Description
keychain, page 162	Enables key chain authentication on a TCP connection between two BGP neighbors.

keychain inheritance-disable

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

keychain inheritance-disable

no keychain inheritance-disable

Syntax Description   This command has no arguments or keywords.

Command Default   Configured key chains for neighbor and session groups are inherited.

Command Modes   Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.

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

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

Task ID

Task ID	Operations
bgp	read, write

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

```
RP/0/0/CPU0:router(config)#router bgp 140
RP/0/0/CPU0:router(config-bgp)# session-group group1
RP/0/0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2
```

```
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# keychain inheritance-disable
```

Related Commands

Command	Description
keychain, page 162	Enables key chain authentication on a TCP connection between two BGP neighbors.

label-allocation-mode

To set the MPLS/VPN label allocation mode, use the **label-allocation-mode** command in VRF configuration mode. To remove the **label-allocation-mode** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

label-allocation-mode [**per-ce** | **per-vrf**]

no label-allocation-mode [**per-ce** | **per-vrf**]

Syntax Description

per-ce	Specifies that the same label is used for all the routes advertised from a unique customer edge (CE) peer or router.
per-vrf	Specifies that the same label is used for all the routes advertised from a unique VRF.

Command Default Per-prefix is the default label allocation mode.

Command Modes VRF configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.8.0	The **per-vrf** keyword was added.

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.

Each prefix that belongs to a VRF instance is advertised with a single label, causing an additional lookup to be performed in the VRF forwarding table to determine the customer edge (CE) next hop for the packet. Use the **label-allocation-mode** command with the **per-ce** keyword to avoid the additional lookup on the PE router and conserve label space. This mode allows the PE router to allocate one label for every immediate next hop. The label is directly mapped to the next hop so there is no VRF route lookup performed during data forwarding. However, the number of labels allocated is one for each CE rather than one for each prefix.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to set the label allocation mode to customer edge:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-ce
```

The following example shows how to set the label allocation mode to VRF:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-vrf
```

local-as

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

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

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

Syntax Description

as-number	Valid autonomous system number.
	Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
	Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
	Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
	Cannot be the autonomous system number to which the neighbor belongs.
no-prepend	(Optional) Specifies that local autonomous system values are not prepended to announcements from the neighbor.
replace-as	(Optional) Specifies that prepend only local autonomous system values to announcements to the neighbor.
dual-as	(Optional) Dual-AS mode.
inheritance-disable	Prevents local AS from being inherited from the parent.

Command Default

The BGP autonomous system number specified in the **router bgp** command is used, except when confederations are in use. The confederation autonomous system is used for external neighbors in an autonomous system that is not part of the confederation.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **no-prepend** and **disable** keywords were added.

Release	Modification
Release 3.3.0	This command was supported in VRF neighbor configuration mode.
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.8.0	The **replace-as** keyword was added.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The **dual-as** keyword was added. The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

You can specify the autonomous system number the local BGP uses to peer with each neighbor. The autonomous system number specified with this command cannot be the local BGP autonomous system number (specified with the **router bgp** command) or the autonomous system number of the neighbor (specified with the **remote-as** command). This command cannot be specified for internal neighbors or for external neighbors in an autonomous system that is part of a confederation.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.

maximum-paths (BGP)

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

maximum-paths {**ebgp**| **ibgp**| **eibgp**} *maximum* [**unequal-cost**]

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

Syntax Description

ebgp	Specifies external BGP multipath peers.
ibgp	Specifies internal BGP multipath peers.
eibgp	Specifies internal and external BGP multipath peers. eiBGP allows simultaneous use of internal and external paths.
maximum	Maximum number of parallel routes that BGP installs in the routing table. Range is 2 to 8
unequal-cost	(Optional) Allows iBGP multipaths to have different BGP next-hop Interior Gateway Protocol (IGP) metrics. This option is available when the **ibgp** keyword is used.

Command Default

One path is installed in the routing table.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	TheThe **ebgp** and **ibgp** keywords were added and the *maximum* range was changed from 1–8 to 2–8.
Release 3.3.0	This command was supported in VRF IPv4 address family configuration mode. The **eibgp** and **unequal-cost** keywords were added.
Release 3.5.0	This command was supported in VRF IPv6 address family configuration mode.

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 **maximum-paths** command to allow the BGP protocol to install multiple paths into the routing table for each prefix. Multiple paths are installed for external peers that are from the same autonomous system and are equal cost (according to the BGP best-path algorithm). Similarly, multiple paths are installed for internal peers that are equal cost based on the BGP best-path algorithm. The IGP metric to the BGP next hop is the same as the best-path IGP metric unless the router is configured for unequal cost iBGP multipath or eiBGP multipath. See *Implementing BGP* in the *Cisco IOS XR Routing Configuration Guide for the Cisco XR 12000 Series Router* for information on the BGP best-path algorithm.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

maximum-prefix (BGP)

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

maximum-prefix *maximum* [*threshold*] [**warning-only**]

no maximum-prefix [*maximum* [*threshold*] [**warning-only**]]

Syntax Description

maximum	Maximum number of prefixes allowed from this neighbor. Range is 1 to 4294967295.
threshold	(Optional) Integer specifying at what percentage of the *maximum* argument value the software starts to generate a warning message. Range is 1 to 100.
warning-only	(Optional) Instructs the software to generate a log message only when the maximum argument value is exceeded, and not terminate the peering.

Command Default

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

IPv4 unicast: 524,288 prefixes

IPv4 multicast: 131,072 prefixes

IPv4 tunnel: 524, 288

IPv6 unicast: 131,072 prefixes

IPv6 multicast: 131,072 prefixes

VPNv4 unicast: 524, 288

VPNv6 unicast: 524, 288

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

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

VPNv4 address family group configuration

IPv4 neighbor address family configuration

IPv6 neighbor address family configuration

VPNv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

IPv4 tunnel neighbor address family configuration

IPv4 tunnel neighbor group address family configuration

IPv4 tunnel address family group configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VPNv4 address family, VPNv4 neighbor address, and VPNv4 neighbor group address family configuration modes.
Release 3.5.0	This command was supported in the following configuration modes: • VPNv6 address family group • VPNv6 neighbor address family • VPNv6 neighbor group address family • IPv4 tunnel address family

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

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

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

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

Task ID

Task ID	Operations
bgp	read, write

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

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

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
clear bgp, page 104	Resets a BGP connection using BGP hard or soft reconfiguration.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.

mpls activate (BGP)

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

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

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

mpls activate *interface id*

no mpls activate *interface id*

Syntax Description	*interface id* Name of the interface.

Command Default   No default behavior or values

Command Modes   Router configuration

Neighbor configuration

IPv4 address family group configuration

VPNv4 address family group configuration

Command History

Release	Modification
Release 3.6.0	This command was introduced.

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

The**mpls activate** command enables MPLS on the interface specified and also adds the implicit null rewrite corresponding to the peer associated with the interface. The interface specified must be the one corresponding to the inter-AS ASBR or CSC peer.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to activate MPLS for InterAS Option B (with confederations):

```
RP/0/0/CPU0:router(config)#router bgp 1

  bgp confederation peers

    2002

    !

  bgp confederation identifier 4589

  bgp router-id 3.3.3.3

  mpls activate

   interface GigabitEthernet0/1/0/0

   !

  address-family ipv4 unicast

   redistribute connected

   !

  address-family vpnv4 unicast

   retain route-target all

   !

  neighbor 10.0.0.9

   remote-as 2002

   address-family ipv4 unicast

    route-policy pass in

    route-policy pass out

    !

   address-family vpnv4 unicast

    route-policy pass in
```

The following example shows how to activate MPLS for CSC (with confederations):

```
router bgp 2002
  bgp confederation peers
   1
   !
  bgp confederation identifier 4589
  bgp router-id 4.4.4.4
  address-family ipv4 unicast
   allocate-label all
   !
  address-family vpnv4 unicast
   retain route-target all
   !
  vrf foo
   rd 1:1
   mpls activate
     interface GigabitEthernet0/1/0/2
```

```
         !
         address-family ipv4 unicast
          redistribute connected
          allocate-label all
         !
         neighbor 10.0.0.1
          remote-as 1
          address-family ipv4 unicast
          !
          address-family ipv4 labeled-unicast
           route-policy pass in
           route-policy pass out
          !
         !
        !
      !
RP/0/0/CPU0:router#show mpls forwarding
  Local  Outgoing    Prefix            Outgoing Next Hop      Bytes
  Label  Label       or ID             Interface
  Switched
  ------ ----------- ----------------- ------------ ---------------
  ------------
  16000  Aggregate   foo: Per-VRF Aggr[V]   \
                                         foo                          0
  16001  Pop         10.0.0.0/16[V]    Gi0/1/0/2    10.0.0.1       44



RP/0/0/CPU0:router#show mpls interfaces
  Interface              LDP      Tunnel   Enabled
  ---------------------- -------- -------- --------
  GigabitEthernet0/1/0/2   No       No       Yes
```

Related Commands

Command	Description
address-family (BGP), page 14	Enters address family configuration mode for configuring BGP routing sessions.

neighbor (BGP)

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

neighbor *ip-address*

no neighbor *ip-address*

Syntax Description

ip-address	IPv4 or IPv6 IP address of the BGP-speaking neighbor.

Command Default   Neighbor mode is not specified.

Command Modes   Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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.

From router configuration mode, you can use this command to enter neighbor configuration mode.

From neighbor configuration mode, you can enter address family configuration for the neighbor by using the **address-family** command, which allows you to configure routing sessions for IP Version 4 and IP Version 6 address prefixes.

The **neighbor** command does not cause the neighbor to be configured and does not result in a peering to be established with the neighbor. To create the neighbor, you configure a remote autonomous system number by entering the **remote-as** command, or the neighbor can inherit a remote autonomous system from a neighbor group or session group if the **use** command is applied.

> **Note**   A neighbor must have must a remote autonomous system number, and an IP address and address family must be enabled on the neighbor.

Unlike IPv4, IPv6 must be enabled before any IPv6 neighbors can be defined. Enable IPv6 in router configuration mode using the **address-family** command.

✎

Note   Configuration for the neighbor cannot occur (peering is not established) until the neighbor is given a remote as-number and neighbor address.

The **no** form of this command causes the peering with the neighbor to be terminated and all configuration that relates to the neighbor to be removed.

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

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

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

Related Commands

Command	Description
address-family (BGP), page 14	Enters address family configuration mode for configuring BGP routing sessions.
remote-as (BGP), page 215	Adds an entry to the BGP neighbor table.
use, page 455	Inherits characteristics from a neighbor group, session group, or address family group.

neighbor-group

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

neighbor-group *name*

no neighbor-group *name*

Syntax Description

name	Neighbor group name.

Command Default   No neighbor group mode is specified.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

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

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

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

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

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

The **no** form of this command ordinarily causes all configuration for the neighbor group to be removed. If using the **no** form would result in a neighbor losing its remote autonomous system number, the configuration

is rejected. In this scenario, the neighbor configuration must be either removed or configured with a remote autonomous system number before the neighbor group configuration can be removed.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to create a neighbor group called group1 that has IP Version 4 (IPv4) unicast and IPv4 multicast activated along with various configuration features. The neighbor group is used by neighbor 10.0.0.1 and neighbor 10.0.0.2, which allows them to inherit the entire group1 configuration.

```
RP/0/0/CPU0:router(config)# router bgp 65530
RP/0/0/CPU0:router(config-bgp)# neighbor-group group1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# remote-as 65535
RP/0/0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 2
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# send-community-ebgp
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# next-hop-self
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/0/CPU0:router(config-bgp)#neighbor 10.0.0.1
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
address-family (BGP), page 14	Enters various address family configuration modes for configuring BGP routing sessions.
neighbor (BGP), page 180	Enters neighbor configuration mode for configuring BGP routing sessions.
use, page 455	Inherits characteristics from a neighbor group, a session group, or an address family group.

network (BGP)

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

network {*ip-address/prefix-length* | *ip-address mask*} [**route-policy** *route-policy-name*]

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

Syntax Description

ip- *address*	Network that BGP advertises.
/ *prefix-length*	Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
ip-address mask	Network mask applied to the *ip-address* argument.
route-policy *route-policy-name*	(Optional) Specifies a route policy to use to modify the attributes of the network.

Command Default   No networks are specified.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **policy** keyword was changed to **route-policy** .
Release 3.3.0	This command was supported in VRF IPv4 address family configuration mode.
Release 3.5.0	This command was supported in VRF IPv6 address family configuration mode.

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.

A network specified with this command is originated and advertised to neighbors only if there exists a route for the network in the routing table. That is, there must be a route learned using local or connected networks, static routing, or a dynamic IGP such as IS-IS or OSPF.

Other than the available system resources on the router, no limit exists on the number of network commands that can be configured.

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

Related Commands

Command	Description
network backdoor, page 186	Specifies a backdoor route to a BGP border router that provides better information about the network.
redistribute (BGP), page 211	Redistributes routes from one routing domain into another routing domain.

network backdoor

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

network {*ip-address/prefix-length* | *ip-address mask*} **backdoor**

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

Syntax Description

ip-address	Network that provides a backdoor route.
/ *prefix-length*	Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
mask	Network mask applied to the *ip-address* argument.

Command Default

No backdoor routes are installed.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF IPv4 address family configuration mode.
Release 3.5.0	This command was supported in VRF IPv6 address family configuration mode.

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.

Configuring the **network backdoor** command does not cause BGP to originate a network, even if an IGP route for the network exists. Ordinarily, the backdoor network would be learned through both an eBGP and

IGP. The BGP best-path selection algorithm does not change when a network is configured as a backdoor network.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows IP Version 4 (IPv4) unicast network 192.168.40.0/24 configured as a backdoor network:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# network 192.168.40.0/24 backdoor
```

Related Commands

Command	Description
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.

next-hop-self

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

next-hop-self [**inheritance-disable**]

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

Syntax Description

inheritance-disable	(Optional) Allows a next-hop calculation override when this feature may be inherited from a neighbor group or address family group.

Command Default

When this command is not specified, the software calculates the next hop for BGP updates accepted by the router.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

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

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command is supported in the following configuration modes: • VPNv4 address family group • VPNv4 neighbor address family • VPNv4 neighbor group address family
Release 3.5.0	This command was supported in the following configuration modes:

Release	Modification
	• VPNv6 address family group
	• VPNv6 neighbor address family
	• VPNv6 neighbor group address family
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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 **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.

Task ID

Task ID	Operations
bgp	read, write

Examples

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

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

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

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

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
use, page 455	Inherits characteristics from a neighbor group, session group, or address family group.

next-hop-unchanged

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

next-hop-unchanged [**inheritance-disable**]

no next-hop-unchanged [**inheritance-disable**]

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

Command Default Overwriting of the next hop is allowed.

Command Modes VPNv4 address family group configuration

VPNv4 neighbor address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

IPv4 labeled-unicast address family configuration

IPv6 labeled-unicast address family configuration

IPv4 address family configuration

IPv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.5.0	This command was supported in the following configuration modes: • VPNv6 address family group • VPNv6 neighbor address family • VPNv6 neighbor group address family
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

Release	Modification
Release 4.0.0	This command was supported in the following address family configuration modes:
	• IPv4 labeled-unicast address family configuration
	• IPv6 labeled-unicast address family configuration
	• IPv4 unicast address family configuration
	• IPv6 unicast address family configuration

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

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

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# af-group group1 address-family vpnv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# next-hop-unchanged disable
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
```

Related Commands

Command	Description
next-hop-self, page 188	Disables next-hop calculation and allows you to insert your own address in the next-hop field of BGP updates.
use, page 455	Inherits characteristics from a neighbor group, session group, or address family group.

nexthop route-policy

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

nexthop route-policy *route-policy-name*

no nexthop route-policy *route-policy-name*

Syntax Description

route-policy-name	Route policy to use for filtering based on next hops.

Command Default   No default behavior or values

Command Modes   IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 3.4.0	This command was introduced.
Release 3.5.0	This command was supported in VPNv6 address family configuration mode.

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 **nexthop route-policy** command to configure route policy filtering using next hops.

The BGP next-hop tracking feature allows you to specify that BGP routes are resolved using only next hops whose routes have the following characteristics:

- To avoid the aggregate routes, the prefix length must be greater than a specified value.

- The source protocol must be from a selected list, ensuring that BGP routes are not used to resolve next hops that could lead to oscillation.

This route policy filtering is possible because RIB identifies the source protocol of a route that resolves a next hop as well as the mask length associated with the route.

The next-hop attach point supports matching using the protocol name and mask length. BGP marks all next hops that are rejected by the route policy as invalid, and no best path is calculated for the routes that use the invalid next hop. The invalid next hops continue to stay in the active cache and can be displayed as part of the **show bgp nexthop** command with an invalid status.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to specify the route policy nexthop_A as the policy to use for filtering next hops:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# nexthop route-policy nexthop_A
```

Related Commands

Command	Description
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.
show bgp nexthops, page 332	Display statistical information about the BGP next hops.

nexthop trigger-delay

To specify the delay for triggering next-hop calculations, use the **nexthop trigger-delay** command in the appropriate configuration mode. To set the trigger delay to the default value, use the **no** form of this command.

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

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

Syntax Description

critical	Specifies critical next-hop events. For example, when the next hop is unreachable.
delay	Trigger delay, in milliseconds. Range is 0 to 4294967295.
non-critical	Specifies noncritical next-hop events. For example, Interior Gateway Protocol (IGP) metric changes.

Command Default

critical : 3000 msec for IPv4 address family and IPv6 address family

critical : msec for VPNv4 address family and VPNv6 address family

non-critical : 10000 msec IPv4, IPv6, VPNv4, and VPNv6 address families

Command Modes

IPv4 address family configuration

Pv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.4.0	This command was changed from **bgp nexthop-trigger-delay** to **nexthop trigger-delay** .
	The supported command mode was changed from Router configuration to the following configuration modes:
	• IPv4 address family configuration
	• IPv6 address family configuration
	• VPNv4 address family configuration
	The **critical** and **non-critical** keywords have been added. The *delay* range has changed from 0 to 300 seconds to 0 to 4294967295 msec.

Release	Modification
Release 3.5.0	This command was supported in VPNv6 address family configuration mode.
Release 3.8.0	The default critical delay value for VPNv4 address family and VPNv6 address family was set to 0 msec.

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 **nexthop trigger-delay** command to allow for a dynamic way for Interior Gateway Protocol (IGP) to converge. This convergence allows BGP to accumulate all notifications and trigger fewer walks, resulting in fewer interprocess communications (IPCs) to the Routing Information Base (RIB) for route addition, deletion, and modification and fewer updates to peers.

Note A high *delay* value can be configured to effectively turn off next-hop tracking.

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

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the critical next-hop trigger delay to 3500 milliseconds:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# nexthop trigger-delay critical 3500
```

nsr (BGP)

To activate Border Gateway Protocol (BGP) nonstop routing (NSR), use the **nsr** command in BGP global configuration mode. To deactivate BGP NSR, use the **no**form of this command.

nsr

no nsr

Syntax Description   This command has no arguments or keywords.

Command Default   BGP NSR is not activated.

Command Modes   Global configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 **nsr** command to enable the Border Gateway Protocol (BGP) Nonstop Routing (NSR) with Stateful Switchover (SSO). This enables all bgp peerings to maintain the BGP state to ensure continuous packet forwarding during events that could interrupt service.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable BGP NSR:

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 120
RP/0/0/CPU0:router(config-bgp)# nsr
```

The following example shows how to disable BGP NSR:

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 120
RP/0/0/CPU0:router(config-bgp)# no nsr
```

Related Commands

Command	Description
router bgp, page 231	Configures the Border Gateway Protocol (BGP) routing process.
nsr process-failures switchover	Configures failover as a recovery action in case of process failures for active instances to switch over to a standby route processor (RP) or a standby distributed route processor (DRP) to maintain nonstop routing (NSR).
show bgp nsr, page 341	Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information.

orf

To specify Outbound Route Filter (ORF) and inbound filtering criteria, use the **orf route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

orf route-policy *route-policy-name*

no orf route-policy *route-policy-name*

Syntax Description

route-policy-name	Name of the route policy.

Command Default

No ORF route policy is defined.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.5.0	This command was supported in VRF IPv6 neighbor address family configuration mode.

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.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to configure outbound and inbound filtering criteria:

```
RP/0/0/CPU0:router(config)#router bgp 6
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbr-af)#orf route-policy policy_A
```

Related Commands

Command	Description
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.

password (BGP)

To enable Message Digest 5 (MD5) authentication on a TCP connection between two Border Gateway Protocol (BGP) neighbors, use the **password** command in an appropriate configuration mode. To disable MD5 authentication, use the **no** form of this command.

password {**clear**| **encrypted**} *password*

no password [**clear** *password*| **encrypted** *password*]

Syntax Description

clear	Specifies that an unencrypted password follows. The password must be a case-sensitive, clear-text unencrypted password.
encrypted	Specifies that an encrypted password follows. The password must be a case-sensitive, encrypted password.
password	Password of up to 80 characters. The password can contain any alphanumeric characters. However, if the first character is a number or the password contains a space, the password must be enclosed in double quotation marks; for example, "2 password."

Command Default

When this command is not specified in the appropriate configuration mode, MD5 authentication is not enabled on a TCP connection between two BGP neighbors.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **0** and **7** keywords were replaced with the **clear** and **encrypted** keywords and the **accept** keyword was removed.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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.

Configure a password to enable authentication between two BGP peers. Use the **password** command to verify each segment sent on the TCP connection between the peers. The same password must be configured on both networking devices, otherwise a connection cannot be made. The authentication feature uses the MD5

algorithm. Specifying this command causes the software to generate and check the MD5 digest on every segment sent on the TCP connection.

Configuring a neighbor password does not cause the existing session for a neighbor to end. However, until the new password is configured on the remote router, the local BGP process does not receive keepalive messages from the remote device. If the password is not updated on the remote device by the end of the hold time, the session ends. The hold time can be changed using the **timers** command or the **timers bgp** command.

If this command is configured for a neighbor group or neighbor address family group, a neighbor using the group inherits the configuration. Values of commands configured specifically for a neighbor overrides inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure neighbor 172.20.1.1 to use MD5 authentication with the password password1:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)#password clear password1
```

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
password-disable, page 203	Overrides any inherited password configuration from a neighbor group or session group for BGP neighbors.
session-group, page 240	Creates a session group and enters session group configuration mode.
timers (BGP), page 446	Set the timers for a specific BGP neighbor.

password-disable

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

password-disable

no password-disable

Syntax Description   This command has no arguments or keywords.

Command Default   Configured passwords for neighbor and session groups are inherited.

Command Modes   Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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.

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

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to disable MD5 authentication for neighbor 172.20.1.1, preventing it from inheriting the password password1 from session group group1:

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

```
RP/0/0/CPU0:router(config-bgp)# session-group group1
RP/0/0/CPU0:router(config-bgp-sngrp)# password clear password1
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# password-disable
```

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
password (BGP), page 201	Enables MD5 authentication on a TCP connection between two BGP neighbors.
session-group, page 240	Creates a session group and enters session group configuration mode.
use, page 455	Inherits characteristics from a neighbor group, a session group, or an address family group.

precedence

To set the precedence level, use the **precedence** command in the appropriate configuration mode. To remove the **precedence** command from the configuration file and restore the system to its default interval values, use the **no** form of this command.

precedence *value*

no precedence [*value*]

Syntax Description	*value* Value of the precedence. The precedence value can be a number from 0 to 7, or it can be one of the following keywords: **critical** —Set packets with critical precedence (5) **flash** — Set packets with flash precedence (3) **flash-override** —Set packets with flash override precedence (4) **immediate** —Set packets with immediate precedence (2) **internet** —Set packets with internetwork control precedence (6) **network** —Set packets with network control precedence (7) **priority** —Set packets with priority precedence (1) **routine** —Set packets with routine precedence (0)

Command Default   No default behavior or values

Command Modes   Neighbor configuration

Neighbor session group configuration

Neighbor group configuration

Command History

Release	Modification
Release 3.4.0	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 **precedence** command to set the precedence value.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the precedence to 2:

```
RP/0/0/CPU0:router(config)# router bgp 5
RP/0/0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/0/CPU0:router(config-bgp-nbr)# precedence 2
```

rd

To configure a route distinguisher, use the **rd** command in VRF configuration mode. To disable the route distinguisher, use the **no** form of this command.

rd {*as-number* : *nn*| *ip-address* : *nn*| **auto**}

no rd {*as-number* : *nn*| *ip-address* : *nn*| **auto**}

Syntax Description

as-number:nn	• *as-number* —16-bit Autonomous system (AS) number of the route distinguisher
	
	   ◦ 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.
	
	 • *nn* —32-bit number

ip-address:nn	IP address of the route distinguisher. • *ip-address* —32-bit IP address • *nn* —16-bit number
auto	Automatically assigns a unique route distinguisher.

Command Default   No default behavior or values

Command Modes   VRF configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.5.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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 **rd** command to make the prefix unique across multiple VRFs.

Auto assignment of route distinguishers can be done only if a router ID is assigned using the **bgp router-id** command in BGP router configuration mode. The unique router ID is used for automatic route distinguisher generation.

The following are restrictions when configuring route distinguishers:

- BGP router-id must be configured before **rd auto** can be configured

- Route distinguisher cannot be changed or removed when an IPv4 unicast address family is configured under VRF.

- BGP router-id cannot be changed or removed when **rd auto** is configured under a VRF.

- When **rd auto** is configured under a VRF, the IP address for the router distinguisher configured under another VRF must be different from that of the BGP router-id

- If a route distinguisher with same IP address as BGP router-id exists, the **rd auto** is not permitted.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to automatically assign a unique route distinguisher to VRF instance vrf-1:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/0/CPU0:router(config-bgp-vrf)# rd auto
```

Related Commands

Command	Description
bgp router-id, page 90	Configures a fixed router ID for a BGP-speaking router.
export route-target, page 154	Configures a VRF export route-target extended community.
import route-target, page 159	Configures a VRF import route-target extended community.

receive-buffer-size

To set the size of the receive buffers for a Border Gateway Protocol (BGP) neighbor, use the **receive-buffer-size** command in an appropriate configuration mode. To remove the **receive-buffer-size** command from the configuration file and restore the system to its default condition in which the software uses the default size, use the **no** form of this command.

receive-buffer-size *socket-size* [*bgp-size*]

no receive-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size, in bytes, of the receive-side socket buffer. Range is 512 to 131072.
bgp-size	(Optional) Size, in bytes, of the receive buffer in BGP. Range is 512 to 131072.

Command Default

socket-size : 32,768 bytes

bgp-size : 4,032 bytes

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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 **receive-buffer-size** command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because it allows the software to process a larger number of packets simultaneously. However, allocating larger buffers consumes more memory on the router.

Note Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the receive buffer sizes for neighbor 172.20.1.1 to be 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# receive-buffer-size 65536 8192
```

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
send-buffer-size, page 234	Sets the size of the send buffers for a BGP neighbor.
session-group, page 240	Creates a session group and enters session group configuration mode.
socket receive-buffer-size, page 435	Sets the size of the receive buffers for all BGP neighbors.

redistribute (BGP)

To redistribute routes from one routing domain into Border Gateway Protocol (BGP), use the **redistribute** command in an appropriate configuration mode. To disable route redistribution, use the **no** form of this command.

Connected

redistribute connected [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute connected [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Enhanced Interior Gateway Routing Protocol (EIGRP)

redistribute eigrp *process-id* [**match** {**external**| **internal**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute eigrp *process-id* [**match** {**external**| **internal**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Intermediate System-to-Intermediate System (IS-IS)

redistribute isis *process-id* [**level**| {**1**| **1-inter-area** | **2**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute isis *process-id* [**level**| {**1**| **1-inter-area** | **2**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Open Shortest Path First (OSPF)

redistribute ospf *process-id*

no redistribute ospf *process-id*

OSPFv3

redistribute ospf *process-id*

no redistribute ospf *process-id*

Routing Information Protocol

redistribute rip [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute rip [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Static

redistribute static [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute static [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Syntax Description

connected	Redistributes connected routes. Connected routes are established automatically when IP is enabled on an interface.

metric *metric-value*	(Optional) Specifies the Multi Exit Discriminator (MED) attribute used for the redistributed route. Range is 0 to 4294967295. Use a value consistent with the destination protocol.
	By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0.
route-policy *route-policy-name*	(Optional) Specifies a configured routing policy to filter redistributed routes. A route policy is used to filter the importation of routes from this source routing protocol to BGP.
eigrp	Specifies that routes are distributed from EIGRP. You must be in IPv4 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode.
process-id	For the **eigrp** keyword, an EIGRP instance name from which routes are to be redistributed.
	For the **isis** keyword, an IS-IS instance name from which routes are to be redistributed.
	For the **ospf** keyword, an OSPF instance name from which routes are to be redistributed.
	The *process-id* value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.
match { **internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**] }	(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following: • **internal** —Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes). • **external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • **nssa-external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed.
isis	Specifies that routes are distributed from the IS-IS protocol.
	Redistribution from IS-IS is allowed under IPv4 unicast, IPv4 multicast, IPv6 unicast, and IPV6 multicast address-families. Redistribution is not allowed under VPNv4 and VPNv6 address-families.
level { **1** | **1-inter-area** | **2** }	(Optional) Specifies the IS-IS level from which routes are redistributed. It can be one of the following: • **1** —Routes are redistributed from Level 1 routes. • **1-inter-area** —Routes are redistributed from Level 1 interarea routes. • **2** —Routes are redistributed from Level 2 routes.

ospf	Specifies that routes are distributed from the OSPF protocol. You must be in IPv4 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode.
ospfv3	Specifies that routes are distributed from the OSPFv3 protocol. You must be in IPv6 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode.
rip	Specifies that routes are distributed from RIP. You must be in IPv4 unicast or multicast address family configuration mode.
static	Redistributes IP static routes.

Command Default

Route redistribution is disabled.

For IS-IS, the default is to redistribute Level 1 and Level 2 routes.

For OSPF, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2.

For OSPFv3, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2

By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0.

metric *metric-value:* 0

match { **internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**]}: If no match is specified, the default is to match all routes.

Command Modes

IPv4 address family configuration, both unicast and multicast (**connected**, **eigrp**, **isis**, **ospf**, **rip**, and **static** are supported)

IPv6 address family configuration, both unicast and multicast (**connected**, **eigrp, isis, ospfv3,** and **static** are supported)

VRF IPv4 address family configuration (**connected** , **eigrp** , **ospf** , **rip** , and **static** are supported)

VRF IPv6 address family configuration (**connected** , **eigrp** , and **static** are supported)

Command History

Release	Modification
Release 3.2	This command was introduced. The **policy** keyword was changed to **route-policy** . The **1-inter-area** and **opsfv3** keywords were added.
Release 3.3.0	The **eigrp** and **rip** keywords were added. This command was supported in VRF IPv4 address family configuration mode.
Release 3.5.0	This command was supported in VRF IPv6 address family configuration mode.

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   When redistributing routes (into BGP) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Each instance of a protocol may be redistributed independently of the others. Changing or removing redistribution for a particular instance does not affect the redistribution capability of other protocols or other instances of the same protocol.

Networks specified using the **network** command are not affected by the **redistribute** command; that is, the routing policy specified in the **network** command takes precedence over the policy specified through the **redistribute** command.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to redistribute IP Version 4 (IPv4) unicast OSPF routes from OSPF instance 110 into BGP:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# redistribute ospf 110
```

Related Commands

Command	Description
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.

remote-as (BGP)

To create a Border Gateway Protocol (BGP) neighbor and begin the exchange of routing information, use the **remote-as** command in an appropriate configuration mode. To delete the entry for the BGP neighbor, use the **no** form of this command.

remote-as *as-number*

no remote-as [*as-number*]

Syntax Description

as-number	Autonomous system (AS) to which the neighbor belongs.
	• Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
	• Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
	• Range for 4-byte Autonomous system numbers (ASNs) is 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 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# session-group group2
RP/0/0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/0/CPU0:router(config-bgp-sngrp)#exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/0/CPU0:router(config-bgp-nbr)#use session-group group2
```

The following example shows how to configure a session group called group2 with an autonomous system number 1. Neighbor 10.0.0.1 is created when it inherits the autonomous system number 1 from session group group2.

```
RP/0/0/CPU0:router(config)#router bgp 1
RP/0/0/CPU0:router(config-bgp)# session-group group2
RP/0/0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group group2
```

Related Commands

Command	Description
neighbor (BGP), page 180	Enters neighbor configuration mode for configuring BGP routing sessions.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
router bgp, page 231	Configures the BGP routing process.
session-group, page 240	Creates a session group and enters session group configuration mode.
use, page 455	Inherits characteristics from a neighbor group, session group, or address family group.

remove-private-as

To remove private autonomous system numbers from autonomous system paths when generating updates to external neighbors, use the **remove-private-as** command in an appropriate configuration mode. To place the router in the default state in which it does not remove private autonomous system numbers, use the **no** form of this command.

remove-private-as [**inheritance-disable**] [**entire-aspath**]

no remove-private-as [**inheritance-disable**] [**entire-aspath**]

Syntax Description

inheritance-disable	(Optional) Permits the feature to be disabled from a neighbor group or address family group instead of being inherited.
entire-aspath	(Optional) Removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private.

Command Default When this command is not specified in the appropriate configuration mode, private autonomous system numbers are not removed from updates sent to external neighbors.

Command Modes IPv4 address family group configuration

IPv6 address family group configuration

VPNv4 address family group configuration

IPv4 neighbor address family configuration

VPNv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the following configuration modes: • VPNv4 address family group

Release	Modification
	• VPNv4 neighbor address family group
	• VRF IPv4 neighbor address family
	• VPNv4 neighbor group address family
Release 3.5.0	This command was supported in the following configuration modes:
	• VPNv6 address family group
	• VPNv6 neighbor address family
	• VRF IPv6 neighbor address family
	• VPNv6 neighbor group address family
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.
Release 3.9.2	The **entire-aspath** keyword was supported.

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.

This feature is available for external BGP (eBGP) neighbors only.

When an update is passed to the external neighbor, the software drops any leading autonomous system sequence in the autonomous system path if the sequence contains only private autonomous system numbers and does not contain the autonomous system number of the neighbor.

If this command is used in a BGP confederation, the element following the confederation portion of the autonomous system path, if a sequence, is considered the leading sequence.

The private autonomous system values range from 64512 to 65535.

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

Use the **entire-aspath** to removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows a configuration that removes the private autonomous system number from the IP Version 4 (IPv4) unicast updates sent to 172.20.1.1:

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

```
RP/0/0/CPU0:router(config-bgp-nbr-af)# remove-private-as
```

The following example shows how to disable the remove private autonomous system number feature for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

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

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
remote-as (BGP), page 215	Allows entries to the BGP neighbor table.

retain local-label

To retain the local label until the network is converged, use the **retain local-label** command in an appropriate address family configuration mode. To disable the retaining of the local label, use the **no** form of this command.

retain local-label *minutes*

no retain local-label

Syntax Description

minutes	Local retention time in minutes. The range is 3 to 60 minutes. The default retention time is 5 minutes.

Command Default

minutes : 5

Command Modes

L2VPN address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 3.9.0	This command was introduced.

Usage Guidelines

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to enable local label retention for 5 minutes:

```
RP/0/0/CPU0:router# configure
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# retain local-label 5
```

Related Commands

Command	Description
additional-paths install backup, page 7	Installs a backup path into the forwarding table
advertise best-external, page 18	Advertises the best-external path to the iBGP and route-reflector peers.

retain route-target

To accept received updates with specified route targets, use the **retain route-target** command in an appropriate configuration mode. To disable the retaining of routes tagged with specified route targets, use the **no** form of this command.

retain route-target {**all**| **route-policy** *route-policy-name*}

no retain route-target [**all**| **route-policy** *route-policy-name*]

Syntax Description

all	Accepts received updates containing at least one route target.
route-policy *router-policy-name*	Accepts received updates accepted by a specified route filter policy.

Command Default

The default is to accept all route targets.

Command Modes

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.5.0	This command was supported in VPNv6 address family configuration mode.

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 **retain route-target** command to configure a route reflector (RR) to retain routes tagged with specific route targets (RT).

A provider edge (PE) router is not required to hold all VPNv4 routes. The PE router holds only routes that match the import RT of the VPNs configured on it, but a RR must retain all VPNv4 routes because it may peer with PE routers and different PEs may require different RT-tagged VPNv4 routes. Configuring an RR to hold only routes that have a defined set of RT communities and configuring some of these RRs to service a different set of VPNs provides scalability to the RRs. A PE can be configured to peer with all RRs that service the VPN routing and forwarding (VRF) instances configured on the PE. When a new VRF is configured with an RT for which the PE does not already hold routes, the PE issues route refresh requests to the RRs and gets the relevant VPN routes.

The **route-policy** *route-policy-name* keyword and argument takes the policy name that lists the extended communities that a path should have for the RR to retain the path.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure RR to retain all routes with the route filter policy ft-policy-A:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# retain route-target route-filter ft-policy-A
```

Related Commands

Command	Description
import route-target, page 159	Configures a VRF import route-target extended community.

route-policy (BGP)

To apply a routing policy to updates advertised to or received from a Border Gateway Protocol (BGP) neighbor, use the **route-policy** command in an appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *route-policy-name* [*parameter1, parameter2, . . . , parametern*] {**in** | **out**}

no route-policy *route-policy-name* [*parameter1, parameter2, . . . , parametern*] {**in** | **out**}

Syntax Description

route-policy-name	Name of route policy. Up to 16 parameters can follow the route-policy-name, enclosed in brackets ([ ]).
in	Applies policy to inbound routes.
out	Applies policy to outbound routes.

Command Default No policy is applied.

Command Modes IPv4 address family group configuration

IPv6 address family group configuration

VPNv4 address family group configuration

IPv4 neighbor address family configuration

VPNv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **policy** keyword was changed to **route-policy** .
Release 3.3.0	This command was supported in the following configuration modes:

Release	Modification
	• VPNv4 address family group
	• VPNv4 neighbor address family
	• VRF IPv4 neighbor address family
	• VPNv4 neighbor group address family
Release 3.5.0	This command was supported in the following configuration modes:
	• VPNv6 address family group
	• VPNv6 neighbor address family
	• VRF IPv6 neighbor address family
	• VPNv6 neighbor group address family
	Up to 16 parameters were supported following the route-policy-name.

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 **route-policy** command to specify a routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes. The **route-policy** command is used to define a policy.

Note   Configuring a large number of uniquely named outbound neighbor policies can adversely affect performance. This is true even if the uniquely named route policies are functionally identical. The user is discouraged from configuring multiple functionally identical route policies for use with this command. For example, if Policy A and Policy B are identical but named for different neighbors, the two policies should be configured as a single policy.

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

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to apply the In-Ipv4 policy to inbound IP Version 4 (IPv4) unicast routes from neighbor 172.20.1.1:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
```

```
RP/0/0/CPU0:router(config-bgp-nbr-af)# route-policy In-Ipv4 in
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
route-policy (RPL)	Defines a route policy and enters route-policy configuration mode.

route-reflector-client

To configure the router as a Border Gateway Protocol (BGP) route reflector and configure the specified neighbor as its client, use the **route-reflector-client** command in an appropriate configuration mode. To disable configuring the neighbor as a client, use the **no** form of this command.

route-reflector-client [**inheritance-disable**]

no route-reflector-client [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Allows the configuration inherited from a neighbor group or address family group to be overridden.

Command Default   The neighbor is not treated as a route reflector client.

Command Modes   IPv4 address family group configuration

IPv6 address family group configuration

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

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the following configuration modes: • VPNv4 address family group • VPNv4 neighbor address family • VPNv4 neighbor group address family
Release 3.5.0	This command was supported in the following configuration modes: • VPNv6 address family group

Release	Modification
	• VPNv6 neighbor address family
	• VPNv6 neighbor group address family
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

This command is restricted to internal BGP (iBGP) neighbors only.

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.

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.

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.

The local router is a route reflector as long as it has at least one route reflector client.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows neighbor at 172.20.1.1 configured as a route reflector client for IP Version 4 (IPv4) unicast routes:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

The following example disables the route-reflector client for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# route-reflector-client
RP/0/0/CPU0:router(config-bgp-afgrp)#exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
```

```
RP/0/0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/0/CPU0:router(config-bgp-nbr-af)# route-reflector-client inheritance-disable
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
bgp cluster-id, page 58	Configures the cluster ID if the BGP cluster has more than one route reflector.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.

router bgp

To configure the Border Gateway Protocol (BGP) routing process, use the **router bgp** command in global configuration 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) is asdot format is 1.0 to 65535.65535.

Command Default   No BGP routing process is enabled.

Command Modes   Global configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.4.0	The *as-number* 4-byte number range 1.0 to 65535.65535 was supported.
Release 3.9.0	Asplain format for 4-byte Autonomous system number notation was supported.

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

Task ID	Operations
rib	read, write

Examples

The following example shows how to configure a BGP process for autonomous system 120:

```
RP/0/0/CPU0:router(config)# router bgp 120
```

selective-vrf-download disable

To disable Selective VRF Download (SVD), use the **selective-vrf-download disable** command in global configuration mode. To enable SVD, use the **no** form of this command.

selective-vrf-download disable

no selective-vrf-download disable

Syntax Description   This command has no arguments or keywords.

Command Default   SVD is enabled.

Command Modes   Global configuration

Command History

Release	Modification
Release 4.1.0	This command was introduced.

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

You must failover the active RP or reload the router after disabling SVD for the configuration change to get activated.

Task ID

Task ID	Operation
ip-services	read, write

Examples   This example shows how to disable selective vrf download:

```
RP/0/0/CPU0:router#configure
RP/0/0/CPU0:router(config)#selective-vrf-download disable
```

Related Commands

Command	Description
show svd role, page 428	Displays Selective VRF Download (SVD) role information.
show svd state, page 430	Displays Selective VRF Download (SVD) state information.

send-buffer-size

To set the size of the send buffers for a Border Gateway Protocol (BGP) neighbor, use the **send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default values, use the **no** form of this command.

send-buffer-size *socket-size* [*bgp-size*]

no send-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size, in bytes, of the send-side socket buffer. Range is 4096 to 131072.
bgp-size	(Optional) Size, in bytes, of the BGP process send buffer. Range is 4096 to 131072.

Command Default

socket-size : 10240 bytes

bgp-size : 4096 bytes

Use the **socket send-buffer-size** command to change the defaults.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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 **send-buffer-size** command to increase the buffer size employed when sending updates to a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on the router.

Note Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses more memory indefinitely.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the send buffer sizes for neighbor 172.20.1.1 to be 8192 bytes for both the socket buffer and the BGP buffer:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# send-buffer-size 8192 8192
```

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
receive-buffer-size, page 209	Sets the size of the receive buffers for a BGP neighbor.
session-group, page 240	Creates a session group and enters session group configuration mode.
socket send-buffer-size, page 437	Sets the size of the send buffers for all BGP neighbors.

send-community-ebgp

To specify that community attributes should be sent to an external Border Gateway Protocol (eBGP) neighbor, use the **send-community-ebgp** command in an appropriate configuration mode. To disable sending community attributes to an eBGP neighbor, use the **no** form of this command.

send-community-ebgp [**inheritance-disable**]

no send-community-ebgp [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Allows configuration inherited from a neighbor group or address family group to be overridden.

Command Default

Community attributes are not sent to eBGP neighbors.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the VRF IPv4 neighbor address family configuration mode.
Release 3.5.0	This command was supported in VRF IPv6 neighbor address family configuration mode.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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 **send-community-ebgp** command to control whether community attributes are sent to eBGP neighbors. It cannot be configured for iBGP neighbors. Communities are always sent to iBGP neighbors.

If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Configuring the command specifically for a neighbor overrides inherited values.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to disable the router that sends community attributes to neighbor 172.20.1.1 for IP Version 4 (IPv4) multicast routes:

```
RP/0/0/CPU0:router(config)#router bgp 140
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# send-community-ebgp
```

The following example shows how to disable the delivery of community attributes to neighbor 172.20.1.1, preventing this feature from being inherited from address family group group1:

```
RP/0/0/CPU0:router(config)#router bgp 140
RP/0/0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-afgrp)# send-community-ebgp
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/0/CPU0:router(config-bgp-nbr-af)# send-community-ebgp inheritance-disable
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
send-extended-community-ebgp, page 238	Specifies that extended community attributes are sent to eBGP neighbors.

send-extended-community-ebgp

To specify that extended community attributes should be sent to external Border Gateway Protocol (eBGP) neighbors, use the **send-extended-community-ebgp** command in an appropriate configuration mode. To disable sending extended community attributes to eBGP neighbors, use the **no** form of this command.

send-extended-community-ebgp [**inheritance-disable**]

no send-extended-community-ebgp [**inheritance-disable**]

Syntax Description

inheritance-disable	(Optional) Allows configurations inherited from a neighbor group or address family group to be overridden.

Command Default   Extended community attributes are not sent to an eBGP neighbor.

Command Modes   IPv4 address family group configuration

IPv6 address family group configuration

IPv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VRF IPv6 neighbor address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the VRF IPv4 neighbor address family configuration mode.
Release 3.5.0	This command was supported in VRF IPv6 neighbor address family configuration mode.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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 **send-extended-community-ebgp** command to control whether extended community attributes are sent to eBGP neighbors. It cannot be used for iBGP neighbors. Extended communities are always sent to iBGP neighbors.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure the router to send extended community attributes to neighbor 172.20.1.1 for IP Version 4 (IPv4) multicast routes:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# send-extended-community-ebgp
```

The following example shows how to disable the delivery of extended community attributes to neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-afgrp)# send-extended-community-ebgp
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/0/CPU0:router(config-bgp-nbr-af)# send-extended-community-ebgp inheritance-disable
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
send-community-ebgp, page 236	Specifies that community attributes should be sent to an eBGP neighbor.

session-group

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

session-group *name*

no session-group *name*

Syntax Description

name	Name of the session group.

Command Default   No session groups are created.

Command Modes   Router configuration

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 **session-group** command to create a session group from which neighbors can inherit configuration that is address family-independent. That is, session groups cannot have address family-specific configuration. This command enters the session group configuration mode in which configuration for a session group is entered.

Many commands can be configured in both session group configuration mode and neighbor configuration mode.

Use of session groups saves time and reduces the router configuration size. Because the configuration of a session group can be inherited by any number of neighbors, use of the group can eliminate the need to copy long or complex configurations on each of a large number of neighbors. A neighbor can inherit all configuration from a session group simply by configuring the **use** command. Specific inherited session group configuration commands can be overridden for a specific neighbor by explicitly configuring the command for the specific neighbor.

The **no** form of this command causes all of the configuration for the session group to be removed. You cannot use the **no** form of this command if removing the group would leave one or more neighbors without a configured remote autonomous system number.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows a session group called group1 that is used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group.

```
RP/0/RSP0RP0/CPU0:router(config)# router bgp 1
RP/0/RSP0RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
```

The following example shows a session group called group1 used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group. However, the **password password1** configuration from group1 is overridden for neighbor 10.0.0.2, using the **password-disable** command in the neighbor 10.0.0.2 configuration submode.

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# session-group group1
RP/0/0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/0/CPU0:router(config-bgp-sngrp)# password password1
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# password-disable
```

session-open-mode

To establish a Border Gateway Protocol (BGP) session with a specific TCP open mode, use the **session-open-mode** command in an appropriate configuration mode. To restore the default state, use the **no** form of this command.

session-open-mode {**active-only**| **both**| **passive-only**}

no session-open-mode [**active-only**| **both**| **passive-only**]

Syntax Description

active-only	Ensures that the BGP session can be established only when the request is initiated by the local end (active-open request) and all passive-open requests (from the other end) are rejected by the local BGP.
both	Allows BGP sessions to be established from both incoming or outgoing TCP connection requests, with one being rejected in the event of a request collision.
passive-only	Ensures that the local BGP does not initiate any TCP open requests and the session can be established only when the request comes from the remote end.

Command Default

The default is **both** .

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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, by default, tries to initiate an active TCP connection whenever a new neighbor is configured. A remote neighbor may also initiate the TCP connection before the local BGP can initiate the connection. This initiation of a TCP connection by a remote neighbor is considered a passive-open request and it is accepted by the local BGP. This default behavior can be modified using the **session-open-mode** command.

Note   The BGP connection is not opened and, as a result the BGP session, is not established if both the peering neighbors use the same nondefault TCP session open mode—active-only or passive-only. If both ends are configured with active-only, each neighbor rejects the TCP open request from the other end. One neighbor must be configured as passive-only or both. Similarly, if both neighbors are configured with passive-only, neither neighbor initiates the TCP open request and the BGP session is not established. Again, one neighbor must be configured as active-only or both. There is one exception. A connection open request from a neighbor that is configured with the TCP session open mode to be passive-only is processed to detect whether there is a connection collision before the request is rejected. This exception enables the local BGP to reset the session if the remote neighbor goes down and it is not detected by the local router.

Use the **session-open-mode** command when it may be necessary to preconfigure a neighbor that does not exist. Ensure that BGP does not spend any time actively trying to set up a TCP session with the neighbor. A BGP session does not come up between two neighbors, both of which configure the same nondefault value (**active-only** or **passive-only** keyword) for this command.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable a BGP session on router bgp 1:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# neighbor 45.67.89.01
RP/0/0/CPU0:router(config-bgp-nbr)# session-open-mode active-only
```

show bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show bgp** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast**| **all**| **tunnel**| **mdt**}| **ipv6** {**unicast**| **multicast**| **all**| **labeled-unicast**}| **all** {**unicast**| **multicast**| **all**| **labeled-unicast**| **mdt**| **tunnel**}| **vpnv4 unicast** [**rd** **rd-address**]| **vrf** {**vrf-name**| **all**} [**ipv4** {**unicast**| **labeled-unicast**}| **ipv6 unicast**]| **vpnv6 unicast** [**rd** **rd-address**]] [**ip-address** [{**mask**| /**prefix-length**} [**longer-prefixes**| **unknown-attributes**| **bestpath-compare**]]] [**standby**] [**detail**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
tunnel	(Optional) Specifies tunnel address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
ip-address	(Optional) Network address, entered to display a particular network in the BGP routing table. If the network address is omitted, then all networks in the

BGP routing table are displayed. If the network mask and prefix length is omitted, then the software displays the longest matching prefix for the network address.

mask	(Optional) Network mask of the BGP route to match.
/ prefix-length	(Optional) Prefix length of the BGP route to match. A slash (/) must precede the decimal value.
longer-prefixes	(Optional) Displays a route with the specified prefix length and more-specific routes if available. The **longer-prefixes** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
unknown-attributes	(Optional) Includes unknown, transitive attributes. The **unknown-attributes** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
bestpath-compare	(Optional) Displays route and best-path comparison information. The **bestpath-compare** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
standby	(Optional) Displays information about the standby card.
detail	(Optional) Displays the prefix details.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **bestpath-compare** keyword was added.
Release 3.3.0	The **vrf** { *vrf-name* | **all** }, **labeled-unicast** , and **vpnv4 unicast** [**rd** *rd-address*] keywords and argument were added.
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families
	The **standby** keyword was removed.

Release	Modification
Release 3.8.0	The 'Last Modified' field was added to show the timestamp when a route was last modified.
	The **standby** keyword was added.
	The **detail** keyword was added to use with the */prefix-length* argument.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
Release 4.0.0	The command output was modified to display BGP add-path information.

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 sessions and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for an address family or a subaddress family, each matching routing table is examined in turn.

Use the **show bgp** *ip-address* { *mask* | */prefix-length* } command to display detailed information for a specific route. If the mask and prefix length are omitted, the details of the longest matching prefix for the IP address are displayed.

Use the **show bgp** command to display all routes in the specified BGP routing table. Use the **show bgp** *ip-address* { *mask* | */prefix-length* } **longer-prefixes** command to display those routes more specific than a particular prefix.

Use the **unknown-attributes** keyword to display details of any transitive attributes associated with a route that are not understood by the local system.

Use the **show bgp** *ip-address/prefix-length* **detail** command to display details of the specified prefix.

Task ID

Task ID	Operations
bgp	read

Examples   The following is sample output from the **show bgp** command in EXEC mode:

```
RP/0/0/CPU0:router#show bgp
 BGP router identifier 172.20.1.1, local AS number 1820
 BGP generic scan interval 60 secs
 BGP table state: Active
 Table ID: 0xe0000000
 BGP main routing table version 3
 Dampening enabled
 BGP scan interval 60 secs

 Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
 Origin codes: i - IGP, e - EGP, ? - incomplete

 Network           Next Hop           Metric LocPrf Weight Path
 * i10.3.0.0/16    172.20.22.1             0    100      0 1800 1239 ?
 *>i               172.20.16.1             0    100      0 1800 1239 ?
 * i10.6.0.0/16    172.20.22.1             0    100      0 1800 690 568 ?
 *>i               172.20.16.1             0    100      0 1800 690 568 ?
 * i10.7.0.0/16    172.20.22.1             0    100      0 1800 701 35 ?
 *>i               172.20.16.1             0    100      0 1800 701 35 ?
 *                 192.168.40.24                         0 1878 704 701 35 ?
 * i10.8.0.0/16    172.20.22.1             0    100      0 1800 690 560 ?
 *>i               172.20.16.1             0    100      0 1800 690 560 ?
 *                 192.168.40.24                         0 1878 704 701 560 ?
 * i10.13.0.0/16   172.20.22.1             0    100      0 1800 690 200 ?
 *>i               172.20.16.1             0    100      0 1800 690 200 ?
 *                 192.168.40.24                         0 1878 704 701 200 ?
 * i10.15.0.0/16   172.20.22.1             0    100      0 1800 174 ?
 *>i               172.20.16.1             0    100      0 1800 174 ?
 * i10.16.0.0/16   172.20.22.1             0    100      0 1800 701 i
 *>i               172.20.16.1             0    100      0 1800 701 i
 *                 192.168.40.24                         0 1878 704 701 i

 Processed 8 prefixes, 8 paths
```

This table describes the significant fields shown in the display.

Table 3: show bgp Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
BGP table state	State of the BGP database.

Field	Description
Table ID	BGP database identifier.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between BGP scans for the specified address family and subaddress family.
Status codes	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): 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. The second character may be (in order of precedence): >—Path is the best path to use for that network. 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. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using 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.
Network	IP prefix and prefix length for a network.

Field	Description
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the path origin code.

The following is sample output from the **show bgp** command with the network specified:

```
RP/0/0/CPU0:router# show bgp 11.0.0.0/24
BGP router table entry for 11.0.0.0/24
 Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 2        2
 Last Modified: Mar  3 16:12:07.147 for 2d21h
 Paths: (3 available, best #1)
   Advertised to update-groups (with more than one peer):
     0.1
   Advertised to peers (in unique update groups):
     10.4.101.1
   Received by speaker 0
   Local
     0.0.0.0 from 0.0.0.0 (10.4.0.1)
       Origin IGP, metric 0, localpref 100, weight 32768, valid, local, best
   Received by speaker 0
   2 3 4
     10.4.101.1 from 10.4.101.1 (10.4.101.1)
       Origin IGP, localpref 100, valid, external
   Received by speaker 0
   Local
     10.4.101.2 from 10.4.101.2 (10.4.101.2)
       Origin IGP, localpref 100, valid, internal
```

This table describes the significant fields shown in the display.

Table 4: show bgp prefix length Field Descriptions

Field	Description
BGP router table entry	Network that is being displayed.
Versions	List of the network versions in each BGP process.

Field	Description
Process	Name of the BGP process.
bRIB/RIB	Version of the network for sending to the RIB. You can compare this version with the bRIB/RIB version for the process (at the top of **show bgp summary**) to verify whether the network has been sent to the RIB.
SendTblVer	Version of the network for advertising to neighbors. This can be compared with the neighbor version to determine whether the network has been advertised to a particular neighbor.
Last Modified	Timestamp when this route was last modified.
Paths	List of paths for the network (that is, routes to reach the network). The number of paths and the index of the best path are given.
not advertised to any peer	Best path was received with a NO_ADVERTISE community and is not advertised to any neighbor.
not advertised to EBGP peer	Best path was received with a NO_EXPORT community and is not advertised to any eBGP neighbor.
not advertised outside local AS	Best path was received with a LOCAL_AS community and is not advertised to peers outside the local AS.
Advertisements of this net are suppressed by an aggregate	Network is a more-specific prefix of a configured aggregate and has been suppressed. It is not advertised to any neighbors unless they have an unsuppress-map configured.
Advertised to update-groups	List of update-groups to which the net has been advertised. Update-groups that have only one peer are not listed here.
Advertised to peers	List of neighbors to which the net has been advertised to. Neighbors that are in one of the update-groups listed above are not listed separately. Only neighbors that are in unique update-groups are listed.
Received by speaker 0	BGP process where the path originated. This is always "speaker 0" for standalone mode. It will be the speaker-id when BGP is in distributed mode.

Field	Description
AS Path	Autonomous system (AS) path that was received for the path. If the AS path is empty, then "Local" is displayed. This is the case for paths that are locally generated on this router or on a neighboring router within the same AS.
aggregated by	If the path is an aggregate, the router-id of the router that performed the aggregation.
suppressed due to dampening	Path has been suppressed due to the configured path dampening.
history entry	Path is withdrawn, but a copy is kept to store the dampening information.
Received from a RR-client	Path was received from a route reflector client.
received-only	If soft reconfiguration inbound is configured, the path was received but dropped by inbound policy, or was accepted and modified. In either event, the received-only value is a copy of the original, unmodified path.
received & used	If soft reconfiguration inbound is configured, the path was received and accepted by inbound policy, but not modified.
stale	Neighbor from which the path was received is down, and the path is kept and marked as stale to support graceful restart.
<nexthop> from <neighbor> (<router-id>)	Next hop for the path. If the next hop is known by a mechanism outside BGP (for example, for redistributed paths), then 0.0.0.0 is displayed. After the next hop, the neighbor from whom the path was received is displayed, along with the neighbor's router-id. If the path was locally generated (for example, an aggregate or redistributed path), then 0.0.0.0 is displayed for the neighbor address.
Origin	IGP: the path originated from an IGP. EGP: the path originated from an EGP. incomplete: the origin of the path is unknown.
metric	MED value of the path.
localpref	Local preference value. This is used to determine the preferred exit point from the local autonomous

Field	Description
	system. It is propagated throughout the local autonomous system.
weight	Locally assigned weight (if not 0) of the path. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
valid	Path is valid and can be considered in the best-path calculation.
redistributed	Path is redistributed through a **redistribute** command.
aggregated	Path is a locally generated aggregate created due to an **aggregate-address** command.
local	Path is a local network source due to a **network** command.
internal	Path was received from an iBGP neighbor.
external	Path was received from an eBGP neighbor.
atomic-aggregate	Path was received with the atomic-aggregate flag set. Some path information has been removed through aggregation.
best	Path is the best path for the network and is used for routing and advertised to peers.
multipath	Path is a multipath and is installed into the RIB along with the best path.
Community	List of communities attached to the path.
Extended community	List of extended communities attached to the path.
Originator	Originator of the path within the AS Cluster list if the path is reflected.
AS Cluster list	List of RR clusters the path has passed through if the path is reflected
Dampinfo	Penalty and reuse information if the path is dampened.
penalty	Current penalty for the path.
flapped	Number of times the path has flapped and the time since the first flap.

Field	Description
reuse in	Time until the path is re-used (undampened).
half life	Configured half-life for the path.
suppress value	Penalty at which the path is suppressed.
reuse value	Penalty at which the path is re-used.
Maximum suppress time	Maximum length of time for which the path can be suppressed.

The following is sample output from the **show bgp** command with the *ip-address/prefix-length* **detail** options:

```
RP/0/0/CPU0:router# show bgp 51.0.0.0/24 detail

Sat Mar 14 00:37:14.109 PST PDT

BGP routing table entry for 51.0.0.0/24

Versions:

  Process           bRIB/RIB   SendTblVer

  Speaker               3           3

    Flags: 0x3e1000, label_retention: not enabled

Last Modified: Mar 13 19:32:17.976 for 05:04:56

Paths: (1 available, best #1)

  Advertised to update-groups (with more than one peer):

    0.3 0.4 0.7 0.8

  Advertised to peers (in unique update groups):

    201.48.20.1

  Path #1: Received by speaker 0

  Flags: 0x1000003

  200 201

    213.0.0.6 from 213.0.0.6 (200.200.3.1)

      Origin IGP, localpref 100, valid, external, best
```

The following is sample output from the show bgp command with the additional paths received from:

```
BGP routing table entry for 51.0.1.0/24, Route Distinguisher: 2:1
Versions:
  Process           bRIB/RIB   SendTblVer
  Speaker              63          63
    Flags: 0x040630f2
Last Modified: Nov 11 12:44:05.811 for 00:00:16
Paths: (3 available, best #2)
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  Path #1: Received by speaker 0
```

```
         Flags: 0x3
         Not advertised to any peer
         111 111 111 111 111 111 111 111
           10.51.0.10 from 10.51.0.10 (11.11.11.11)
             Origin IGP, metric 0, localpref 100, valid, external
             Received Path ID 0, Local Path ID 0, version 0
             Extended community: RT:55:1
         Path #2: Received by speaker 0
         Flags: 0x5060007
         Advertised to CE peers (in unique update groups):
           10.51.0.10
         561 562 563 564 565
           13.0.6.50 from 13.0.6.50 (13.0.6.50)
             Received Label 16
             Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported
             Received Path ID 0, Local Path ID 1, version 63
             Extended community: RT:55:1
         Path #3: Received by speaker 0
         Flags: 0x4060007
         Not advertised to any peer
         591 592 593 594 595
           13.0.9.50 from 13.0.9.50 (13.0.9.50)
             Received Label 16
             Origin IGP, localpref 100, valid, internal, backup, add-path, import-candidate,
imported
             Received Path ID 0, Local Path ID 4, version 63
             Extended community: RT:22:232 RT:55:1
```

This is sample output to explain 'import suspect' state and 'import-suspect' field in **show bgp** command output:

```
RP/0/0/CPU0:router#show bgp vpnv4 unicast rd 11:111 100.16.11.0/24
BGP routing table entry for 100.16.11.0/24, Route Distinguisher: 11:111
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker           1834195     1834195
Paths: (2 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Path #1: Received by speaker 0
  11
    1:16.16.16.16 (metric 30) from 55.55.55.55 (16.16.16.16)
      Received Label 19602
     Origin incomplete, localpref 100, valid, internal, best, import-candidate, not-in-vrf,
 import suspect
      Extended community: RT:11:11
      Originator: 16.16.16.16, Cluster list: 55.55.55.55
  Path #2: Received by speaker 0
  11
    1:16.16.16.16 (metric 30) from 88.88.88.88 (16.16.16.16)
      Received Label 19602
      Origin incomplete, localpref 100, valid, internal, not-in-vrf, import suspect
      Extended community: RT:11:11
      Originator: 16.16.16.16, Cluster list: 88.88.88.88
```

The **show bgp** command output displays 'import suspect' when potential import oscillation has been detected for the prefix. Import of such a prefix is not affected. However, import of the prefix can be dampened in future if the oscillation continues. If the oscillation stops during the next import run, the prefix will no longer be marked 'import supect'.

Related Commands

Command	Description
aggregate-address, page 24	Creates an aggregate entry in a BGP routing table.
bgp default local-preference, page 67	Changes the default local preference value.

Command	Description
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default subaddress Family Identifier (SAFI) for the current session.
show bgp cidr-only, page 272	Displays routes with nonnatural netmasks.
show bgp community, page 276	Displays routes belonging to the specified communities.
show bgp inconsistent-as, page 297	Displays networks with inconsistent origin autonomous system.
show bgp regexp, page 380	Displays routes matching an AS path regular expression.
show bgp route-policy, page 385	Displays networks that match a route policy.
show bgp summary, page 396	Displays the status of all BGP connections.
show bgp truncated-communities, page 409	Displays networks with community lists truncated by policy.

show bgp advertised

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

show bgp [**ipv4** { **all**| **labeled-unicast**| **mdt**| **multicast**| **tunnel**| **unicast**}] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**ipv6** { **all**| **labeled-unicast**| **multicast**| **unicast**}] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**all** { **all** | **labeled-unicast** | **multicast**| **tunnel** | **unicast** }] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **advertised** [**neighbor** *ip-address*] [**standby**] [**summary**]

show bgp [**vrf** {*vrf-name*| **all**} [**ipv4** | {**labeled-unicast** | **unicast**}| **ipv6 unicast**]] **advertised** [**neighbor** *ip-address*] [**standby**] [**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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 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.
ip-address	(Optional) IP address of the neighbor.
standby	(Optional) Displays information about the standby card.
summary	(Optional) Displays a summary of advertisements.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The following keywords and arguments were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast**}] • [**vpnv4 unicast**] [**rd** *rd-address* | **vrf** *vrf-name*]
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added. The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that is configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp advertised** command to display the routes that have been advertised to peers or a specific peer. To preview advertisements that would be sent to a peer under a particular policy, even if the corresponding update messages have not been generated yet, use the **show bgp policy** command.

Note   When you issue the **show bgp advertised** command, a route is not displayed in the output unless an advertisement for that route has already been sent (and not withdrawn). If an advertisement for the route has not yet been sent, the route is not displayed.

Use the **summary** keyword to display a summary of the advertised routes. If you do not specify the **summary** keyword, the software displays detailed information about the advertised routes.

Note   The **show bgp advertised** command does not display the application of any outbound policy in the route details it displays. Consequently, this command provides only an indication of whether a particular route has been advertised, rather than details of which attributes were advertised. Use the **show bgp policy sent-advertisements** command to display the attributes that are advertised.

Task ID

Task ID	Operations
bgp	read

Examples   The following is sample output from the **show bgp advertised** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp advertised neighbor 10.0.101.4 summary

Network     Next Hop          From           AS Path
1.1.1.0/24  10.0.101.1        10.0.101.1     2 3 222 333 444 555 i
1.1.2.0/24  10.0.101.1        10.0.101.1     3 4 5 6 7 i
1.1.3.0/24  10.0.101.1        10.0.101.1     77 88 33 44 55 99 99 99 i
```

```
1.1.4.0/24  10.0.101.1          10.0.101.1   2 5 6 7 8 i
1.1.7.0/24  10.0.101.1          10.0.101.1   3 5 i
1.1.8.0/24  10.0.101.1          10.0.101.1   77 88 99 99 99 i
```

This table describes the significant fields shown in the display.

Table 5: 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/0/0/CPU0:router# show bgp advertised neighbor 172.72.77.1

172.16.0.0/24 is advertised to 172.72.77.1
  Path info:
    neighbor: Local          neighbor router id: 172.74.84.1
    valid  redistributed  best
  Attributes after inbound policy was applied:
 next hop: 0.0.0.0
    MET ORG AS
    origin: incomplete  metric: 0
    aspath:
10.52.0.0/16 is advertised to 172.72.77.1
  Path info:
    neighbor: Local Aggregate  neighbor router id: 172.74.84.1
    valid  aggregated  best
  Attributes after inbound policy was applied:
 next hop: 0.0.0.0
    ORG AGG ATOM
    origin: IGP  aggregator: 172.74.84.1 (1)
    aspath:
```

This table describes the significant fields shown in the display.

Table 6: 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.
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.

Field	Description
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 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:

Field	Description
	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.
community	Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities:
	Local-AS—Community with value 4294901812. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.
	no-advertise—Community with value 4294901813. Routes with this community value are not advertised to any BGP peers.
	no-export—Community with value 4294901814. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation with the local router.
Extended community	Extended community attributes associated with the path. For known extended community types, the following codes may be displayed:
	RT—Route target community
	SoO—Site of Origin community

Field	Description
	LB—Link Bandwidth community
Originator	Router ID of the originating router when route reflection is used.
Cluster lists	Router ID or cluster ID of all route reflectors through which the route has passed.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default subaddress Family Identifier (SAFI) for the current session.
route-policy (BGP), page 225	Applies a route policy to incoming and outgoing routes.
rd, page 207	Filters routes using a prefix list.
show bgp policy, page 348	Displays information about BGP advertisements under a proposed policy.
show bgp policy, page 348 **sent-advertisements**	Previews advertisements to peers, including details of advertised attributes.

show bgp af-group

To display information about Border Gateway Protocol (BGP) configuration for address family groups, use the **show bgp af-group** command in EXEC mode.

show bgp af-group *group-name* {**configuration** [**defaults**] [**nvgen**]| **inheritance**| **users**}

Syntax Description

group-name	Name of the address family group to display.
configuration	(Optional) Displays the effective configuration for the af-group, including any settings that have been inherited from af-groups used by this af-group.
defaults	(Optional) Displays all configuration settings, including any default settings.
nvgen	(Optional) Displays output in the format of **show running-config** output. If the **defaults** keyword is also specified, the output is not suitable for cutting and pasting into a configuration session.
inheritance	Displays the af-groups from which this af-group inherits configuration settings.
users	Displays the neighbors, neighbor groups, and af-groups that inherit configuration from this af-group.

Command Default

No default behavior or value

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes 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 af-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of an af-group, taking into account any configuration that may be inherited from other af-groups through the **use af-group** command. The source of each command is shown.

If the **defaults** keyword is specified, all configuration for the af-group, including default values, is shown. Default configuration is identified in the show output. Use the **nvgen** keyword to display configuration formatted in the style of the **show running-config** command. This output is suitable for cutting and pasting into configuration sessions.

Use the **show bgp af-group** command with the *group-name* **inheritance** argument and keyword to display the address family groups from which the specified af-group inherits configuration.

Use the **show bgp af-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and af-groups that inherit configuration from the specified af-group.

Task ID

Task ID	Operations
bgp	read

Examples

The following af-group configuration is used in the examples:

```
af-group group3 address-family ipv4 unicast
remove-private-AS
soft-reconfiguration inbound
!
af-group group1 address-family ipv4 unicast
use af-group group2
maximum-prefix 2500 75 warning-only
default-originate
soft-reconfiguration inbound disable
!
af-group group2 address-family ipv4 unicast
use af-group group3
send-community-ebgp
send-extended-community-ebgp
capability orf prefix both
```

The following is sample output from the **show bgp af-group** command with the **configuration** keyword in EXEC mode. The source of each command is shown in the right column. For example, **default-originate** is configured directly on **af-group group1** , and the **remove-private-AS** command is inherited from af-group group2, which in turn inherits it from af-group group3.

```
RP/0/0/CPU0:router# show bgp af-group group1 configuration

af-group group1 address-family ipv4 unicast
  capability orf prefix both              [a:group2]
  default-originate                       []
  maximum-prefix 2500 75 warning-only     []
  remove-private-AS                       [a:group2 a:group3]
  send-community                          [a:group2]
  send-extended-community                 [a:group2
```

The following is sample output from the **show bgp af-group** command with the **users** keyword:

```
RP/0/0/CPU0:router# show bgp af-group group2 users

IPv4 Unicast: a:group1
```

The following is sample output from the **show bgp af-group** command with the **inheritance** keyword. This example shows that the specified af-group group1 directly uses the group2 af-group, which in turn uses the group3 af-group:

```
RP/0/RSP0RP0/CPU0:router# show bgp af-group group1 inheritance

IPv4 Unicast: a:group2 a:group3
```

Table 7: show bgp af-group Field Descriptions, page 266 describes the significant fields shown in the display.

This table describes the significant fields shown in the display.

Table 7: show bgp af-group Field Descriptions

Field	Description
[ ]	Configures the command directly on the specified address family group.
a:	Indicates the name that follows is an address family group.
n:	Indicates the name that follows is a neighbor group.
[dflt]	Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the **defaults** keyword is specified.
<not set>	Indicates that the configuration is disabled by default. This field may be shown when the **defaults** keyword is specified.

Related Commands

Command	Description
af-group, page 22	Configures a BGP address family group.
show bgp neighbors, page 310	Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups.
show bgp neighbor-group, page 306	Displays information about configuration for neighbor groups.
use, page 455 **af-group**	Configures an af-group to inherit the configuration of a specified af-group.

show bgp attribute-key

To display all existing attribute keys, use the **show bgp attribute-key**command in EXEC mode.

show bgp {ipv4| ipv6| all| vpnv4 unicast| vrf| vpnv6 unicast} attribute-key [standby]

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.
mdt		(Optional) Specifies multicast distribution tree (MDT) address prefixes.
all		(Optional) For subaddress family, specifies prefixes for all subaddress families.
ipv6		(Optional) Specifies IP Version 6 address prefixes.
vpnv4 unicast		(Optional) Specifies VPNv4 unicast address families.
vrf		(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name		(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.
ipv6 unicast		(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast		(Optional) Specifies VPNv6 unicast address families.
standby		(Optional) Displays information about the standby card.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The **ipv4** { **unicast** | **labeled-unicast** } keyword was added.
Release 3.4.0	The following keywords and argument were added: • **vpnv4 unicast** • **vrf** (*vrf-name* | **all** } The **labeled-unicast** keyword was supported. The **count-only** keyword was removed.
Release 3.5.0	The **vpnv6 unicast** keywords were added. The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp attribute-key**command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp all all attribute-key
```

```
Address Family: IPv4 Unicast
============================

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 109
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network           Next Hop        AttrKey
*> 1.1.0.0/16        0.0.0.0         0x00000002
*> 10.0.0.0/16       0.0.0.0         0x00000002
*> 12.21.0.0/16      0.0.0.0         0x00000002
*> 194.3.192.1/32    10.0.101.1      0x00000009
*> 194.3.192.2/32    10.0.101.1      0x00000009
*> 194.3.192.3/32    10.0.101.1      0x00000009
*> 194.3.192.4/32    10.0.101.1      0x00000009
*> 194.3.192.5/32    10.0.101.1      0x00000009

Processed 8 prefixes, 8 paths

Address Family: IPv4 Multicast
==============================

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 15
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network           Next Hop        AttrKey
*> 194.3.193.2/32    10.0.101.1      0x00000009
*> 194.3.193.3/32    10.0.101.1      0x00000009

Processed 2 prefixes, 2 paths

Address Family: IPv6 Unicast
============================

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 19
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network           Next Hop        AttrKey
*> 2222::1111/128    2222::2         0x00000009
*> 2222::1112/128    2222::2         0x00000009

Processed 2 prefixes, 2 paths
```

This table describes the significant fields shown in the display.

Table 8: show bgp attribute-key Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.

Field	Description
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
BGP scan interval	Interval (in seconds) between scans.
Status codes	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):
	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.
	The second character may be (in order of precedence):
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:
	i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a **network** or **aggregate-address** command.
	e—Entry 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.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.

Field	Description
AttrKey	Key associated with the route attribute.
Processed *n* prefixes, *n* paths	Number of prefixes and number of paths processed for the table.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

show bgp cidr-only

To display routes with nonnatural network masks, also known as classless interdomain routing (CIDR) routes, use the **show bgp cidr-only** command in EXEC mode.

show bgp [**ipv4**| **vrf**] **cidr-only** [**standby**]

Syntax Description

ipv4	(Optional) Specifies the IP Version 4 address family.
unicast	(Optional) Specifies the unicast address family.
multicast	(Optional) Specifies the multicast address family.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) For subaddress family, specifies all subaddress families.
tunnel	(Optional) Specifies the tunnel address family.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(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.
standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used. This command is applicable only for IPv4 prefixes. If the default address family is not IPv4, then the **ipv4** keyword must be used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **count-only** keyword was added.
Release 3.3.0	The following keywords and arguments were added:

Release	Modification
	• **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }] • [**vpnv4 unicast**] [**rd** *rd-address*]
Release 3.4.0	The **labeled-unicast** keyword was supported. The **count-only** keyword was removed.
Release 3.5.0	The **tunnel** and **mdt** keywords were supported under the **ipv4** address family. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for subaddress family, all subaddress family routing tables are examined.

The **show bgp cidr-only** command applies only for IPv4 prefixes. If the **ipv4** keyword is not specified and the default address family is not IPv4, the command is not available.

Use the **show bgp cidr-only** command to display CIDR routes. Routes that have their correct class (class A, B, or C) prefix length are not displayed.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp cidr-only** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp cidr-only

BGP router identifier 172.20.1.1, local AS number 1820
```

```
BGP main routing table version 2589
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop        Metric  LocPrf  Weight  Path
*> 192.0.0.0/8     192.168.72.24   0       1878    ?
*> 192.168.0.0/16  192.168.72.30   0       108     ?
```

This table describes the significant fields shown in the display.

Table 9: show bgp cidr-only Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	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): 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. The second character may be (in order of precedence): >—Path is the best path to use for that network. 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. The third character may be: i—Path was learned by an internal BGP (iBGP) session.

Field	Description
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a **network** or **aggregate-address** command. e—Entry 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.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

Related Commands

Command	Description
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, page 244	Displays BGP routes.

show bgp community

To display routes that have the specified Border Gateway Protocol (BGP) communities, use the **show bgp community** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast** | **all**| **tunnel**| **mdt**}] **community** *community-list* [**exact-match**]

show bgp [**ipv6** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **community** *community-list* [**exact-match**]

show bgp [**all** {**unicast**| **multicast**| **labeled-unicast** | **all**| **tunnel**}] **community** *community-list* [**exact-match**]

show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **community** *community-list* [**exact-match**]

show bgp [**vrf** {*vrf-name*| **all**} [**ipv4** | {**unicast**| **labeled-unicast**}| **ipv6 unicast**]] **community** *community-list* [**exact-match**]

show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **community** *community-list* [**exact-match**]

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.

vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
community	Specifies that only routes with communities specified by *community-list* is displayed.
community-list	Between one and seven communities. Each community can be a number in the range from 1 to 4294967295, a community specified in AA:NN format, or one of the following well-known communities:
	local-AS —Well-known community with value 4294901812. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.
	no-advertise —Well-known community with value 4294901813. Routes with this community value are not advertised to any BGP peers.
	no-export —Well-known community with value 4294901814. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.
	internet —Well-known community whose value is not defined in BGP RFC. IOS XR BGP uses a value of 0 for the internet community. Routes with this community are advertised to all peers without any restrictions.
	For the AA:NN format:
	AA—Range is 0 to 65535.
	NN—Range is 1 to 4294967295.
	Up to seven community numbers can be specified.
exact-match	(Optional) Displays those routes that have communities exactly matching the specified communities.

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **count-only** keyword was added.
Release 3.3.0	The following keywords and arguments were added:
	• **vrf** { *vrf-name* | **all** }
	• [**ipv4** { **unicast** | **labeled-unicast** }]
	• [**vpnv4 unicast**] [**rd** *rd-address*]

Release	Modification
Release 3.4.0	The **labeled-unicast** keyword was supported.
	The **count-only** keyword was removed.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.
	The **standby** keyword was removed.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or the subaddress family, each matching routing table is examined in turn.

If more than seven communities are required, it is necessary to configure a route policy and use the show bgp route-policy, page 385 command.

Use the **exact-match** keyword to display only those routes with a set of communities exactly matching the list of specified communities. If you omit the **exact-match** keyword, those routes containing at least the specified communities are displayed.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp community** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp community 1820:1 exact-match

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 55
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*  10.13.0.0/16     192.168.40.24                        0 1878 704 701 200 ?
*  10.16.0.0/16     192.168.40.24                        0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 10: show bgp community Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	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): 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. The second character may be (in order of precedence): >—Path is the best path to use for that network. 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. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:

Field	Description
	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.
Network	IP prefix and prefix length for a network.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

Related Commands

Command	Description
aggregate-address, page 24	Creates an aggregate entry in a BGP routing table.
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

Command	Description
show bgp, page 244	Displays BGP routes.

show bgp convergence

To display whether a specific address family has reached convergence, use the **show bgp convergence** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast** | **all**| **tunnel**| **mdt**}] **convergence**

show bgp [**ipv6** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **convergence**

show bgp [**all** {**unicast**| **multicast**| **labeled-unicast** | **all**| **mdt**| **tunnel**}] **convergence**

show bgp [**vpnv4 unicast**] **convergence**

show bgp [**vpnv4 unicast**] **convergence**

Syntax Description

ipv4	(Optional) Specifies the IP Version 4 address family.
unicast	(Optional) Specifies the unicast address family.
multicast	(Optional) Specifies the multicast address family.
labeled-unicast	(Optional) Specifies unicast address prefixes.
all	(Optional) For subaddress family, specifies all subaddress families.
tunnel	(Optional) Specifies tunnel address prefixes.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies the IP Version 6 address family.
all	(Optional) For address family, specifies all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 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

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The **vpnv4 unicast** and **labeled-unicast** keywords were added.

Release	Modification
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** keywords were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.
	The **standby** keyword was removed.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp convergence** command to see if there is any pending work for BGP to perform. The software checks the following conditions to determine whether the specified address family has converged. If all the conditions are true, the address family is considered converged.

- All received updates have been processed and best routes selected.

- All selected routes have been installed in the global Routing Information Base (RIB).

- All selected routes have been advertised to peers, including any peers that are not established (unless those peers have been administratively shut down). See the **shutdown (BGP)** command for more information about administrative shutdown.

While testing that all selected routes have been advertised to peers, the **show bgp convergence** command checks the size of the write queue for each neighbor. Because this queue is shared by all address families, there is a small possibility that the command indicates the address family has not converged when, in fact, it has converged. This could happen if the neighbor write queue contained messages from some other address family.

If the specified address family has not converged, the **show bgp convergence** command output does not indicate the amount of work that is pending. To display this information, use the **show bgp summary** command.

Task ID

Task ID	Operations
bgp	read

Examples

The following shows the result of using the **show bgp convergence** command for an address family that has converged:

```
RP/0/0/CPU0:router# show bgp convergence

Converged.
All received routes in RIB, all neighbors updated.
All neighbors have empty write queues.
```

The following shows the result of using the **show bgp convergence** command for an address family that has not converged:

```
RP/0/0/CPU0:router# show bgp convergence

Not converged.
Received routes may not be entered in RIB.
One or more neighbors may need updating.
```

This table describes the significant fields shown in the display.

Table 11: show bgp convergence Field Descriptions

Field	Description
Converged/Not converged	Specifies whether or not all routes have been installed in the RIB and updates have been generated and sent to all neighbors.
[All] Received routes...	For convergence, all routes must have been installed into the RIB and all updates must have been generated. For non-convergence, some routes may not be installed in the RIB, or some routes that have been withdrawn have not yet been removed from the RIB, or some routes that are up to date in the RIB have not been advertised to all neighbors.
[All | One or more] neighbors...	Specifies the status of neighbor updating.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

Command	Description
show bgp summary, page 396	Displays the status of BGP peer connections.
shutdown (BGP), page 431	Disables a neighbor without removing all of its configuration.

show bgp dampened-paths

To display Border Gateway Protocol (BGP) dampened routes, use the **show bgp dampened-paths** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **dampened-paths** [**standby**]

show bgp [**ipv6** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **dampened-paths** [**standby**]

show bgp [**all** {**unicast**| **multicast**| **labeled-unicast** | **all**| **tunnel**}] **dampened-paths** [**standby**]

show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **dampened-paths** [**standby**]

show bgp [**vrf** {*vrf-name*| **all**} [**ipv4** | {**unicast**| **labeled-unicast**}| **ipv6 unicast**]] **dampened-paths** [**standby**]

show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **dampened-paths** [**standby**]

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.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
standby	(Optional) Displays information about the standby card.

Command Default　　If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes　　EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The following keywords and arguments were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }] • [**vpnv4 unicast**] [**rd** *rd-address*]
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp dampened-paths**command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp dampened-paths

BGP router identifier 10.2.0.1, local AS number 3
BGP generic scan interval 60 secs
BGP main routing table version 7
Dampening enabled
BGP scan interval 60 secs
Status codes:s suppressed, d damped, h history, * valid, > best
             i - internal, S stale

Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          From            Reuse    Path
*d 10.0.0.0        10.0.101.35     00:01:20 35 i
```

This table describes the significant fields shown in the display.

Table 12: show bgp dampened-paths Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	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): 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.

Field	Description
	*—Path is valid.
	The second character may be (in order of precedence):
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:
	i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** 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.
Network	IP prefix and prefix length for a network.
From	Neighbor from which the route was received.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

Related Commands

Command	Description
aggregate-address, page 24	Creates an aggregate entry in a BGP routing table.
bgp dampening, page 64	Enables BGP route dampening or changes various BGP route dampening factors.
clear bgp dampening, page 109	Clears BGP route dampening information and unsuppresses the suppressed routes.

Command	Description
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp flap-statistics, page 291	Displays BGP routes that have flapped.
show bgp neighbors, page 310	Displays information about the TCP and BGP connections to neighbors.

show bgp flap-statistics

To display information about Border Gateway Protocol (BGP) paths that have flapped, use the **show bgp flap-statistics** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **flap-statistics** [**regexp** *regular-expression*|**route-policy** *route-policy-name* |**cidr-only** |{*ip-address*| {*mask*| */prefix-length*}}] [**longer-prefixes**] [**detail**] [**standby**]

show bgp [**ipv6** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **flap-statistics** [**regexp** *regular-expression*|**route-policy** *route-policy-name* |**cidr-only** |{*ip-address*| {*mask*| */prefix-length*}}] [**longer-prefixes**] [**detail**] [**standby**]

show bgp [**all** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **flap-statistics** [**regexp** *regular-expression*|**route-policy** *route-policy-name* |**cidr-only** |{*ip-address*| {*mask*| */prefix-length*}}] [**longer-prefixes**] [**detail**] [**standby**]

show bgp [**vpnv4 unicast** [**rd** *rd-address*]] **flap-statistics** [**regexp** *regular-expression*|**route-policy** *route-policy-name* |**cidr-only** |{*ip-address*| {*mask*| */prefix-length*}}] [**longer-prefixes**] [**detail**] [**standby**]

show bgp [**vrf** {*vrf-name*| **all**} [**ipv4** | {**unicast**| **labeled-unicast**}| **ipv6 unicast**]] **flap-statistics** [**regexp** *regular-expression*|**route-policy** *route-policy-name* |**cidr-only** |{*ip-address*| {*mask*| */prefix-length*}}] [**longer-prefixes**] [**detail**] [**standby**]

show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **flap-statistics** [**regexp** *regular-expression*|**route-policy** *route-policy-name* |**cidr-only** |{*ip-address*| {*mask*| */prefix-length*}}] [**longer-prefixes**] [**detail**] [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
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.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
regexp *regular-expression*	(Optional) Displays flap statistics for all paths that match the regular expression.
route-policy *route-policy-name*	(Optional) Displays flap statistics for a route policy.
cidr-only	(Optional) Displays only routes whose prefix length does not match the classful prefix length for that network. The **cidr-only** keyword can be specified only if the address family is IPv4.
ip-address	(Optional) Flap statistics for a network address only.
mask	(Optional) Network mask applied to the *ip-address* argument.
/ *prefix-length*	(Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value.
longer-prefixes	(Optional) Displays flap statistics for the specified prefix and more-specific prefixes. The **longer-prefixes** keyword is available when the *ip-address* and *mask* or */prefix-length* arguments are specified.
detail	(Optional) Displays dampening parameters for the path. The **detail** keyword cannot be specified if the **longer-prefixes** keyword is specified. The **detail** keyword is available when the *ip-address* argument or *ip-address* and *mask* or */prefix-length* arguments are specified.
standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **filter-list** *access-list* keyword and argument were removed.
	The **route-policy** *route-policy-name* keyword and argument were added.

Release	Modification
Release 3.3.0	The following keywords and arguments were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }] • [**vpnv4 unicast**] [**rd** *rd-address*]
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Flap statistics are maintained only for paths if dampening is enabled using the **bgp dampening** command. If dampening is not enabled, the **show bgp flap-statistics** command does not display any paths.

If no arguments or keywords are specified, the software displays flap statistics for all paths for the specified address family. You can use the **regexp** , **filter-list**, **cidr-only**, and **longer-prefixes** options to limit the set of paths displayed.

If you specify a network address without a mask or prefix length, the longest matching prefix for the network address is displayed. When displaying flap statistics for a single route, use the **detail** keyword to display dampening parameters for the route.

Task ID

Task ID	Operations
bgp	read

Examples   The following is sample output from the **show bgp flap-statistics** command:

```
RP/0/0/CPU0:router# show bgp flap-statistics

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 26180
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network            From          Flaps Duration Reuse    Path
*d 10.0.0.0        172.20.16.177  4     00:13:31 00:18:10 100
*d 10.10.0.0       172.20.16.177  4     00:02:45 00:28:20 100
```

The following is sample output from the **show bgp flap-statistics** command with the **detail** keyword in EXEC mode:

```
RP/0/0/CPU0:router# show bgp flap-statistics 172.31.12.166 detail

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 738
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
  Network           From            Flaps Duration Reuse     Path
 h 172.31.12.166    10.0.101.1      6     00:03:28           2 2000 3000

   Half life      Suppress       Reuse penalty    Max. supp. time
   00:15:00       2000           750              01:00:00
```

This table describes the significant fields shown in the display.

Table 13: show bgp flap-statistics Field Descriptions

Field	Description
BGP route identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening has been enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	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):

Field	Description
	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.
	The second character may be (in order of precedence):
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:
	i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network 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.
Network	IP prefix and prefix length for a network that is dampened.
From	IP address of the peer that advertised this route.
Flaps	Number of times the route has flapped.
Duration	Time (in hours:minutes:seconds) since the first flap.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path of the route that is being dampened.

Field	Description
Half life	Half-life value used when dampening this route. The half-life is the amount of time that must elapse to reduce the reuse penalty by half. The half-life value is specified using the **bgp dampening** command.
Suppress	Suppress value used to dampen this route. The suppress value is the value that the penalty must exceed for the route to be suppressed. The suppress value can be configured using the **bgp dampening** command.
Reuse penalty	Reuse penalty used to dampen this route. The penalty must fall below the reuse penalty for the route to be unsuppressed. The reuse penalty can be configured using the **bgp dampening** command.
Max supp. time	Maximum length of time that the route may be suppressed due to dampening. The maximum suppress time can be configured using the **bgp dampening** command.

Related Commands

Command	Description
bgp dampening, page 64	Enables BGP route dampening or changes various BGP route dampening factors.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp dampened-paths, page 286	Displays the BGP dampened routes.
show bgp neighbors, page 310	Displays information about BGP neighbors.

show bgp inconsistent-as

To display Border Gateway Protocol (BGP) routes originated from more than one autonomous system, use the **show bgp inconsistent-as** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast** | **all**| **tunnel**| **mdt**}] **inconsistent-as** [**standby**]

show bgp [**ipv6** {**unicast**| **multicast**| **labeled-unicast** | **all**}] **inconsistent-as** [**standby**]

show bgp [**all** {**unicast**| **multicast**| **labeled-unicast** | **all**| **tunnel**| **mdt**}] **inconsistent-as** [**standby**]

show bgp vpnv4 unicast [**rd** *rd-address*] **inconsistent-as** [**standby**]

show bgp [**vrf** {*vrf-name*| **all**} [**ipv4** | {**unicast**| **labeled-unicast**}| **ipv6 unicast**]] **inconsistent-as** [**standby**]

show bgp [**vpnv6 unicast** [**rd** *rd-address*]] **inconsistent-as** [**standby**]

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.

standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The following keywords and arguments were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }] • [**vpnv4 unicast**] [**rd** *rd-address*]
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added. The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword

is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

Use the **show bgp inconsistent-as**command to search through all prefixes in the specified BGP routing table and display the paths for any prefix that has inconsistent originating autonomous system numbers. The originating autonomous system is the last autonomous system number displayed in the path field and should be the same for all paths.

If a prefix has one or more paths originating from different autonomous systems, all paths for that prefix are displayed.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp inconsistent-as** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp inconsistent-as

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 1129
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop        Metric          LocPrf Weight Path
*  10.0.0.0      172.16.232.55        0           0 300 88 90 99 ?
*>               172.16.232.52     2222           0 400 ?
*  172.16.0.0    172.16.232.55        0           0 300 90 99 88 200 ?
*>               172.16.232.52     2222           0 400 ?
* 192.168.199.0  172.16.232.55        0           0 300 88 90 99 ?
*>               172.16.232.52     2222           0 400 ?
```

This table describes the significant fields shown in the display.

Table 14: show bgp inconsistent-as Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in

Field	Description
	the table. The first character may be (in order of precedence):
	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.
	The second character may be (in order of precedence):
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:
	i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** 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.
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

Field	Description
	system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

Related Commands

Command	Description
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default -safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

show bgp labels

To display Border Gateway Protocol (BGP) routes and their incoming and outgoing labels, use the **show bgp labels**command in EXEC mode.

show bgp labels

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.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
rd *rd-address*	(Optional) Displays routes with a specific route distinguisher.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.4.0	This command was introduced.
Release 3.5.0	The following keywords were added: • **vpnv6 unicast** • **ipv6** { **unicast** | **labeled-unicast** } The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp labels**command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp vrf BAR ipv4 unicast labels

BGP VRF BAR, state: Active BGP Route Distinguisher: 100:1 BGP router identifier 10.1.1.1,
local AS number 100 BGP table state: Active BGP main routing table version 12

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop        Rcvd Label      Local Label
Route Distinguisher: 100:1 (default for vrf BAR)
*> 20.1.1.1/32      10.0.101.1      16              nolabel
*> 20.1.1.2/32      10.0.101.1      16              nolabel
*> 20.1.1.3/32      10.0.101.1      16              nolabel
*> 20.1.1.4/32      10.0.101.1      16              nolabel
*> 20.1.1.5/32      10.0.101.1      16              nolabel

Processed 5 prefixes, 5 paths
```

This table describes the significant fields shown in the display.

Table 15: show bgp labels Field Descriptions

Field	Description
BGP Route Distinguisher	BGP route distinguisher.
BGP router identifier	BGP identifier for the local system.

Field	Description
local AS number	Autonomous system number for the local system.
BGP table state	State of the BGP routing table.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Status codes	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):
	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.
	The second character may be (in order of precedence):
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:
	i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using 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.
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

Field	Description
	entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Rcvd Label	Received label.
Local Label	Local label.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default subaddress Family Identifier (SAFI) for the current session.

show bgp neighbor-group

To display information about the Border Gateway Protocol (BGP) configuration for neighbor groups, use the **show bgp neighbor-group** command in EXEC mode.

show bgp neighbor-group *group-name* {**configuration [defaults] [nvgen]**| **inheritance**| **users**}

Syntax Description

group-name	Name of the address family group to display.
configuration	(Optional) Displays the effective configuration for the neighbor group, including any configuration inherited by this neighbor group.
defaults	(Optional) Displays all configuration, including default configuration.
nvgen	(Optional) Displays output in **show running-config** command output. If the **defaults** keyword is also specified, the output is not suitable for cutting and pasting into a configuration session.
inheritance	Displays the af-groups, session groups, and neighbor groups from which this neighbor group inherits configuration.
users	Displays the neighbors and neighbor groups that inherit configuration from this neighbor group.

Command Default   No default behavior or value

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines   To use this command, you must be in a user group associated with a task group that includes 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 neighbor-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a neighbor group, including any configuration inherited from session groups, address family groups, and neighbor groups through application of the **use** command. The source of each configured command is also displayed.

Use the **defaults** keyword to display all configuration for the neighbor group, including default configuration. The command output identifies default onfiguration. Use the **nvgen** keyword to display configuration in the

output form of **show running-config**command. Output in this form is suitable for cutting and pasting into a configuration session.

The**show bgp neighbor-group** command with the *group-name* **inheritance** argument and keyword displays the session groups, address family groups, and neighbor groups from which the specified neighbor group inherits configuration.

The **show bgp neighbor-group** *group-name* command displays the neighbors and neighbor groups that inherit configuration from the specified neighbor group.

Task ID

Task ID	Operations
bgp	read

Examples

The examples use the following configuration:

```
af-group group3 address-family ipv4 unicast
 remove-private-AS
 soft-reconfiguration inbound
!
af-group group2 address-family ipv4 unicast
 use af-group group3
 send-community-ebgp
 send-extended-community-ebgp
 capability orf prefix both
!
session-group group3
 dmzlink-bw
!
neighbor-group group3
 use session-group group3
 timers 30 90
!
neighbor-group group1
 remote-as 1982
 use neighbor-group group2
 address-family ipv4 unicast
 !
!
neighbor-group group2
 use neighbor-group group3
 address-family ipv4 unicast
  use af-group group2
  weight 100
 !
```

The following is sample output from the **show bgp neighbor-group** command with the **configuration** keyword:

```
RP/0/0/CPU0:router# show bgp neighbor-group group1 configuration

neighbor-group group1
 remote-as 1982                    []
 timers 30 90                      [n:group2 n:group3]
 dmzlink-bw                        [n:group2 n:group3 s:group3]
 address-family ipv4 unicast       []
  capability orf prefix both       [n:group2 a:group2]
  remove-private-AS                [n:group2 a:group2 a:group3]
  send-community-ebgp              [n:group2 a:group2]
  send-extended-community-ebgp     [n:group2 a:group2]
  soft-reconfiguration inbound     [n:group2 a:group2 a:group3]
  weight 100                       [n:group2]
```

The configuration source is shown to the right of each command. In the output, the **remote-as** command is configured directly on neighbor group group1, and the **send-community-ebgp** command is inherited from neighbor group group2, which in turn inherits the setting from af-group group2.

The following is sample output from the **show bgp neighbor-group** command with the **users** keyword. This output shows that the group1 neighbor group inherits session (address family-independent configuration parameters) from the group2 neighbor group. The group1 neighbor group also inherits IPv4 unicast configuration parameters from the group2 neighbor group:

```
RP/0/0/CPU0:router# show bgp neighbor-group group2 users

Session:      n:group1
IPv4 Unicast: n:group1
```

The following is sample output from the **show bgp neighbor-group** command with the **inheritance** keyword. This output shows that the specified neighbor group group1 inherits session (address family-independent configuration) from neighbor group group2, which inherits its own session from neighbor group group3. Neighbor group group3 inherited its session from session group group3. It also shows that the group1 neighbor-group inherits IPv4 unicast configuration parameters from the group2 neighbor group, which in turn inherits them from the group2 af-group, which itself inherits them from the group3 af-group:

```
RP/0/0/CPU0:router# show bgp neighbor-group group1 inheritance

Session:      n:group2 n:group3 s:group3
IPv4 Unicast: n:group2 a:group2 a:group3
```

This table describes the significant fields shown in the display.

Table 16: show bgp neighbor-group Field Descriptions

Field	Description
[ ]	Configures the command directly on the specified address family group.
s:	Indicates the name that follows is a session group.
a:	Indicates the name that follows is an address family group.
n:	Indicates the name that follows is a neighbor group.
[dflt]	Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the **defaults** keyword is specified.
<not set>	Indicates that the default is for the setting to be disabled. This field may be shown when the **defaults** keyword is specified.

Related Commands

Command	Description
af-group, page 22	Configures a BGP address family group.
session-group, page 240	Creates a session group and enters session group configuration mode.
show bgp af-group, page 264	Displays information about configuration for address family groups.
show bgp neighbors, page 310	Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups.
show bgp session-group, page 390	Displays information about the BGP configuration for session groups.
show running-config	Displays the contents of the currently running configuration or a subset of that configuration.
use, page 455	Inherits configuration from a neighbor group, a session group, or an address family group.

show bgp neighbors

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

show bgp neighbors [**performance-statistics**| **missing-eor**] [**standby**]

show bgp neighbors *ip-address*[**advertised-routes**| **dampened-routes**| **flap-statistics**| **performance-statistics**| **received** | {**prefix-filter**| **routes**}| **routes**] [**standby**]

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

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 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 receive an end-of-record (EOR) in read-only mode.
ip-address	(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.
standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced. The **received routes** keyword was added.
Release 3.3.0	The following keywords and argument were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }]

Release	Modification
	• **vpnv4 unicast**
	• **missing-eor**
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.
	The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.
Release 3.9.0	Asplain format for 4-byte autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
Release 4.0.0	Show output was modified to include information on BGP additional paths send and receive information.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify which routing table should be examined. If the **all** keyword is specified for address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp neighbors** command to display detailed information about all neighbors or a specific neighbor. Use the **performance-statistics** keyword to display information about the work related to specific neighbors done by the BGP process.

Use the **show bgp neighbors** command with the *ip-address* **received prefix-filter** argument and keyword to display the Outbound Route Filter (ORF) received from a neighbor.

Use the **advertised-routes** keyword to display a summary of the routes advertised to the specified neighbor.

Use the **dampened-routes** keyword to display routes received from the specified neighbor that have been suppressed due to dampening. For more details, see the **show bgp dampened-paths** command.

To display information about flapping routes received from a neighbor, use the **flap-statistics** keyword. For more details, see the **show bgp flap-statistics** command.

To display the routes received from a neighbor, use the **routes** keyword. For more details, see the **show bgp** command.

Use the **show bgp neighbor** command with the *ip-address* **configuration** argument and keyword to display the effective configuration of a neighbor, including configuration inherited from session groups, neighbor groups, or af-groups through application of the **use** command. Use the **defaults** keyword to display the value of all configurations for the neighbor, including default configuration. Use the **nvgen** keyword to display configuration output format of the **show running-config** command. Output in this format is suitable for cutting and pasting into a configuration session. Use the **show bgp neighbors** command with the *ip-address* **inheritance** argument and keyword to display the session groups, neighbor groups, and af-groups from which the specified neighbor inherits configuration.

Task ID

Task ID	Operations
bgp	read

Examples

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

```
RP/0/0/CPU0:router# show bgp neighbors 10.0.101.1

BGP neighbor is 10.0.101.1, remote AS 2, local AS 1, external link
 Description: routem neighbor
 Remote router ID 10.0.101.1
  BGP state = Established, up for 00:00:56
  TCP open mode: passive only

  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 17: 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	• 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	Autonomous system number of the local system. • 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.

Field	Description
Administratively shut down	Neighbor connection is disabled using the **shutdown** command.
remote router ID	Router ID (an IP address) of the neighbor.
Neighbor under common administration	Neighbor is internal or a confederation peer.
BGP state	Internal state of this BGP connection.
BFD enabled	Status of bidirectional forwarding detection.
TCP open mode	TCP mode used in establishing the BGP session. The following valid TCP mode are supported: • default—Accept active/passive connections • passive-only—Accept only passive connections • active-only—Accept only active connections initiated by the router
Last read	Time since BGP last read a message from this neighbor.
hold time	Hold time (in seconds) used on the connection with this neighbor.
keepalive interval	Interval for sending keepalives to this neighbor.
DMZ-link bandwidth	DMZ link bandwidth for this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. The following valid BGP capabilities are supported: • 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,

Field	Description
	the neighbor also supports the displayed address family.
Received	Number of messages received from this neighbor, the number of notification messages received and processed from this neighbor, and the number of messages that have been received, but not yet processed.
Sent	Number of messages sent to this neighbor, the number of notification messages generated to be sent to this neighbor, and the number of messages queued to be sent to this neighbor.
Minimum time between advertisement runs	Advertisement interval (in seconds) for this neighbor.
For Address Family	Information that follows is specific to the displayed address family.
BGP neighbor version	Last version of the BGP database that was sent to the neighbor for the specified address family.
Update group	Update group to which the neighbor belongs.
Route reflector client	Indicates that the local system is acting as a route reflector for this neighbor.
Inbound soft reconfiguration allowed	Indicates that soft reconfiguration is enabled for routes received from this neighbor. **Note** If the neighbor has route refresh capability, then soft configuration received-only routes are not stored by the local system unless "override route refresh" is displayed.
eBGP neighbor with no inbound or outbound policy: defaults to drop	Indicates that the neighbor does not have an inbound or outbound policy configured using the **route-policy (BGP)** command. Hence, no routes are accepted from or advertised to this neighbor.
Private AS number removed from updates to this neighbor	Indicates that remove-private-AS is configured on the specified address family for this neighbor.
NEXT_HOP is always this router	Indicates that next-hop-self is configured on the specified address family for this neighbor.
Community attribute sent to this neighbor	Indicates that send-community-ebgp is configured on the specified address family for this neighbor.

Field	Description
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: • 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.

Field	Description
Route refresh request	Number of route refresh requests sent and received from this neighbor.
Outbound Route Filter (ORF)	"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. **Note** A received outbound route filter may be displayed using the **show bgp neighbors** command with the **received prefix-filter** keywords.
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	Indicates whether the condition map selects routes that should be advertised, or routes that should not be advertised: Exist—Routes advertised if permitted by the condition route map. Non-exist—Routes advertised if denied by the condition route map.
accepted prefixes	Number of prefixes accepted.

Field	Description
Prefix advertised	Number of prefixes advertised to the neighbor during the lifetime of the current connection with the neighbor.
suppressed	Number of prefix updates that were suppressed because no transitive attributes changed from one best path to the next. **Note** Update suppression occurs only for external BGP neighbors.
withdrawn	Number of prefixes withdrawn from the neighbor during the lifetime of the current connection with the neighbor.
maximum limit	Maximum number of prefixes that may be received from the neighbor. If "(warning-only)" is displayed, a warning message is generated when the limit is exceeded, otherwise the neighbor connection is shut down when the limit is exceeded.
Threshold for warning message	Percentage of maximum prefix limit for the neighbor at which a warning message is generated.
Connections established	Number of times the router has established a BGP peering session with the neighbor.
dropped	Number of times that a good connection has failed or been taken down.
Last reset due to	Reason that the connection with the neighbor was last reset.
Time since last notification sent to neighbor	Amount of time since a notification message was last sent to the neighbor.
Error Code	Type of notification that was sent. The notification data, if any, is also displayed.
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.

Field	Description
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/0/0/CPU0:router# show bgp neighbors 172.20.16.178 routes

BGP router identifier 172.20.16.181, local AS number 1
BGP main routing table version 27
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network           Next Hop              Metric LocPrf Weight Path
*> 10.0.0.0       172.20.16.178         40             0 10 ?
*> 10.22.0.0      172.20.16.178         40             0 10 ?
```

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

```
RP/0/0/CPU0:router# show bgp neighbors 10.0.101.1 dampened-routes

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          From             Reuse    Path
*d 10.0.0.0         10.0.101.1       00:59:30 2 100 1000 i
*d 11.0.0.0         10.0.101.1       00:59:30 2 100 1000 i
*d 12.0.0.0         10.0.101.1       00:59:30 2 100 1000 i
*d 13.0.0.0         10.0.101.1       00:59:30 2 100 1000 i
*d 14.0.0.0         10.0.101.1       00:59:30 2 100 1000 i
```

This table describes the significant fields shown in the display.

Table 18: show bgp neighbors routes Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in

Field	Description
	the table. The first character may be (in order of precedence): 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. The second character may be (in order of precedence): >—Path is the best path to use for that network. 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. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** 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.
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

Field	Description
	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/0/0/CPU0:router# show bgp neighbors 10.0.101.1 flap-statistics

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          From           Flaps Duration Reuse    Path
 h 10.1.0.0         10.0.101.1      5008  2d02h             2 5000 1000
 h 10.2.0.0         10.0.101.1      5008  2d02h             2 2000 3000
 h 10.2.0.0         10.0.101.1      5008  2d02h             2 9000 6000
*d 10.0.0.0         10.0.101.1      5008  2d02h    00:59:30 2 100 1000
 h 10.0.0.0/16      10.0.101.1      5008  2d02h             2 100 102
*d 10.11.0.0        10.0.101.1      5008  2d02h    00:59:30 2 100 1000
*d 10.12.0.0        10.0.101.1      5008  2d02h    00:59:30 2 100 1000
*d 10.13.0.0        10.0.101.1      5008  2d02h    00:59:30 2 100 1000
*d 10.14.0.0        10.0.101.1      5008  2d02h    00:59:30 2 100 1000
 h 192.168.0.0/16   10.0.101.1      5008  2d02h             2 100 101
```

This table describes the significant fields shown in the display.

Table 19: 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.
Status codes	Status of the table entry. The status is displayed as a three-character field at the beginning of each line in

Field	Description
	the table. The first character may be (in order of precedence):
	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.
	The second character may be (in order of precedence):
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:
	i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** 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.
Network	IP prefix and prefix length for a network.
From	Neighbor from which the route was received.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

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

```
RP/0/0/CPU0:router# show bgp neighbors 10.0.101.2 performance-statistics
```

```
BGP neighbor is 10.0.101.2, remote AS 1
  Read 3023 messages (58639 bytes) in 3019 calls (time spent: 1.312 secs)
  Read throttled 0 times
  Processed 3023 inbound messages (time spent: 0.198 secs)
  Wrote 58410 bytes in 6062 calls (time spent: 3.041 secs)
  Processing write list: wrote 0 messages in 0 calls (time spent: 0.000 secs)
  Processing write queue: wrote 3040 messages in 3040 calls (time spent: 0.055 secs)

  Received 3023 messages, 0 notifications, 0 in queue
  Sent 3040 messages, 0 notifications, 0 in queue
```

This table describes the significant fields shown in the display.

Table 20: 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	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): 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. The second character may be (in order of precedence): 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. The third character may be: i—Path was learned by an internal BGP (iBGP) session.

Field	Description
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP prefix and prefix length for a network.
From	IP address of the peer that advertised this route.
Flaps	Number of times the route has flapped.
Duration	Time (in hours:minutes:seconds) since the router noticed the first flap.
Reuse	Time (in hours:minutes:seconds) after which the path is made available.
Path	Autonomous system path to reach the destination network.

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

```
RP/0/0/CPU0:router# show bgp neighbors 10.0.101.1 configuration

neighbor 10.0.101.1
 remote-as 2                   []
 bfd fast-detect               []
 address-family ipv4 unicast   []
  policy pass-all in           []
  policy pass-all out          []
 address-family ipv4 multicast []
  policy pass-all in           []
  policy pass-all out          []
```

This table describes the significant fields shown in the display.

Table 21: show bgp neighbors performance-statistics Field Descriptions

Field	Description
Read	Indicates the number of messages received from the neighbor, the total size of received messages, the number of read operations performed, and the real time spent (in seconds) by the process performing read operations for this neighbor.

Field	Description
Read throttled	Number of times that reading from the TCP connection to this neighbor has been throttled. Throttling is due to a backlog of messages that have been read but not processed.
inbound messages	Number of read messages that have been processed, and the real time spent processing inbound messages for this neighbor.
Wrote	Amount of data that has been sent to this neighbor, number of write operations performed, and the real time spent by the process performing write operations for this neighbor.
Processing write list	Number of messages written from the write list to this neighbor, number of times the write list has been processed, and real time spent processing the write list. **Note** Write lists typically contain only update messages.
Processing write queue	Number of messages written from the write queue to this neighbor, number of times the write queue has been processed, and real time spent processing the write queue.
Received	Number of messages received from this neighbor, number of notification messages received and processed from this neighbor, and number of messages that have been received, but not yet processed.
Sent	Number of messages sent to this neighbor, number of notification messages generated to be sent to this neighbor, and number of messages queued to be sent to this neighbor.

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

```
RP/0/0/CPU0:router# show bgp neighbors 10.0.101.1 configuration

neighbor 10.0.101.1
 remote-as 2                 []
 bfd fast-detect             []
 address-family ipv4 unicast   []
  policy pass-all in          []
  policy pass-all out         []
 address-family ipv4 multicast []
  policy pass-all in          []
  policy pass-all out         []
```

This table describes the significant fields shown in the display.

Table 22: 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.

The following sample output shows sample output from **show bgp neighbors** command with additional paths send and receive capabilities advertised to neighbors:

```
BGP neighbor is 80.0.0.30
 Remote AS 100, local AS 100, internal link
 Remote router ID 33.33.33.33
  BGP state = Established, up for 19:54:12
  NSR State: None
  Last read 00:00:25, Last read before reset 19:54:54
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:02, attempted 19, written 19
  Second last write 00:01:02, attempted 19, written 19
  Last write before reset 19:54:54, attempted 29, written 29
  Second last write before reset 19:54:59, attempted 19, written 19
  Last write pulse rcvd  Nov 11 12:58:03.838 last full not set pulse count 2407
  Last write pulse rcvd before reset 19:54:54
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 19:54:54, second last 19:54:54
  Last KA expiry before reset 00:00:00, second last 00:00:00
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 19:54:54, second last 19:54:59
  Precedence: internet
  Non-stop routing is enabled
  Graceful restart is enabled
  Restart time is 120 seconds
  Stale path timeout time is 360 seconds
  Neighbor capabilities:            Adv       Rcvd
    Route refresh:                  Yes       Yes
    4-byte AS:                      Yes       Yes
    Address family IPv4 Unicast:    Yes       Yes
    Address family IPv4 Labeled-unicast: Yes        Yes
    Address family VPNv4 Unicast:   Yes       Yes
    Address family IPv6 Unicast:    Yes       Yes
    Address family VPNv6 Unicast:   Yes       Yes
    Address family IPv4 MDT:        Yes       Yes
  Message stats:
    InQ depth: 0, OutQ depth: 0
                    Last_Sent              Sent  Last_Rcvd              Rcvd
    Open:           Nov 10 17:03:52.731       2 Nov 10 17:03:52.730       2
    Notification:   ---                       0 ---                       0
    Update:         Nov 10 17:05:02.435      20 Nov 10 17:04:58.812      12
    Keepalive:      Nov 11 12:58:03.632    1197 Nov 11 12:57:40.458    1196
    Route_Refresh:  ---                       0 ---                       0
    Total:                                 1219                        1210
  Minimum time between advertisement runs is 0 secs
```

```
For Address Family: IPv4 Unicast
 BGP neighbor version 13
 Update group: 0.9
 NEXT_HOP is always this router
 AF-dependant capabilities:
   Graceful Restart capability advertised and received
     Neighbor preserved the forwarding state during latest restart
     Local restart time is 120, RIB purge time is 600 seconds
     Maximum stalepath time is 360 seconds
     Remote Restart time is 120 seconds
   Additional-paths Send: advertised and received
   Additional-paths Receive: advertised and received
 Route refresh request: received 0, sent 0
 0 accepted prefixes, 0 are bestpaths
 Prefix advertised 10, suppressed 0, withdrawn 0, maximum limit 524288
 Threshold for warning message 75%
 AIGP is enabled
 An EoR was received during read-only mode
 Last ack version 13, Last synced ack version 0
 Outstanding version objects: current 0, max 1
 Additional-paths operation: Send and Receive

For Address Family: IPv4 Labeled-unicast
 BGP neighbor version 13
 Update group: 0.4 (Update Generation Throttled)

 AF-dependant capabilities:
   Graceful Restart capability advertised and received
     Neighbor preserved the forwarding state during latest restart
     Local restart time is 120, RIB purge time is 600 seconds
     Maximum stalepath time is 360 seconds
     Remote Restart time is 120 seconds
   Additional-paths Send: received
   Additional-paths Receive: received
 Route refresh request: received 0, sent 0
 0 accepted prefixes, 0 are bestpaths
 Prefix advertised 2, suppressed 0, withdrawn 0, maximum limit 131072
 Threshold for warning message 75%
 AIGP is enabled
 An EoR was received during read-only mode
 Last ack version 13, Last synced ack version 0
 Outstanding version objects: current 0, max 1
 Additional-paths operation: None
```

Related Commands

Command	Description
clear bgp, page 104	Resets a BGP connection or session.
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, page 244	Displays entries in the BGP routing table.

Command	Description
show bgp dampened-paths, page 286	Displays BGP dampened routes.
show bgp flap-statistics, page 291	Displays BGP routes that have flapped.
show bgp neighbor-group, page 306	Displays information about the BGP configuration for neighbor groups.
shutdown (BGP), page 431	Disables a neighbor without removing all of its configuration.

show bgp neighbors nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information across neighbors, use the **show bgp neighbors nsr** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **all**}| **ipv6** {**unicast**| **multicast**| **all**}| **vpnv4 unicast**| **vpnv6 unicast**| **vrf** {**all**| *vrf_name*}] **neighbors nsr** [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
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.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
vrf	(Optional) Specifies VPN routing and forwarding (VRF) instance.
vrf_name	(Optional) Name of a VRF.
all	(Optional) For VRF, specifies all VRFs.
standby	(Optional) Displays information about the standby card.

Command Default   No default behavior or values.

Command Modes   EXEC

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp neighbors nsr** command with the **standby** keyword:

```
RP/0/0/CPU0:router# show bgp neighbors nsr standby

BGP neighbor is 2.2.2.2
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0


BGP neighbor is 10.0.101.5
  BGP state = Established, up for 05:19:00
  NSR state = NSR Ready
  Outstanding Postits: 0


BGP neighbor is 10.1.0.5
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0
```

This table describes the significant fields shown in the display.

Table 23: show bgp neighbors nsr Field Descriptions

Field	Description
BGP state	Displays BGP neighbor peering state.
NSR state	Displays BGP neighbor NSR state.
Outstanding Postits	Displays the postit counters of pending events.

Related Commands

Command	Description
nsr (BGP), page 197	Activates the Border Gateway Protocol (BGP) nonstop routing (NSR).
show bgp summary nsr, page 402	Displays the Border Gateway Protocol (BGP) nonstop routing (NSR) information.
show bgp summary, page 396	Displays the status of all Border Gateway Protocol (BGP) connections.

show bgp nexthops

To display statistical information about the Border Gateway Protocol (BGP) next hops, use the **show bgp nexthops** command in EXEC mode.

show bgp nexthops [**statistics**] [**speaker** *speaker-id*] [**standby**]

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
statistics	(Optional) Specifies nexthop statistics.
speaker *speaker-id*	(Optional) Specifies a speaker process ID.
standby	(Optional) Displays information about the standby card.

Command Default   No default behavior or value

Command Modes   EXEC

Command History

Release	Modification
Release 3.4.0	This command was introduced.
Release 3.5.0	The following keywords were added: • **vpnv6 unicast** • **statistics** The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The RefCount value was changed to address family/all format.
Release 3.8.0	The **standby** keyword was added.

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 **show bgp nexthops** command displays statistical information about next-hop notifications, the time spent processing the notifications, and details about each next-hop that has been registered with the Routing Information Base (RIB).

Use the **vrf** *vrf-name* keyword and argument to display only the next-hops present in the specified VPN routing and forwarding (VRF) instance.

The next-hop information is displayed for all active speaker processes in distributed mode. Each speaker displays a set of next-hops that belongs to the prefixes received by the speaker and next hops that belong to best paths that were received by other speaker processes. Use the **speaker** *speaker-id* keyword and argument to display information for only the specified speaker process. The distributed mode must be defined using the **distributed speaker** command for the **speaker** keyword to be available.

Task ID

Task ID	Operations
bgp	read

Examples   The following is sample output from the **show bgp nexthops** command with the VRF specified:

```
RP/0/0/CPU0:router# show bgp vrf all nexthops

Fri Mar 13 17:05:40.656 UTC

VRF: 900
```

```
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000001
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric          Notf        LastRIBEvent RefCount
10.0.101.201    [UR]            4294967295      0/0          1d22h (Reg)       0/3
90.0.0.2        [R][C][NL]               0      1/0          1d22h (Cri)     20/23


VRF: 901
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000002
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric          Notf        LastRIBEvent RefCount
10.0.101.201    [UR]            4294967295      0/0          1d22h (Reg)       0/3
91.0.0.2        [R][C][NL]               0      1/0          1d22h (Cri)     10/13


VRF: 902
========

Total Nexthop Processing
```

```
    Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000003
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric          Notf      LastRIBEvent RefCount
10.0.101.201    [UR]            4294967295       0/0       1d22h (Reg)       0/3
92.0.0.2        [R][C][NL]              0         1/0       1d22h (Cri)      10/13


VRF: 903
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000004
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric          Notf      LastRIBEvent RefCount
10.0.101.201    [UR]            4294967295       0/0       1d22h (Reg)       0/3
93.0.0.2        [R][C][NL]              0         1/0       1d22h (Cri)      10/13


VRF: 904
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
```

```
      Received: 82y48w
      Bestpaths Deleted: 0
      Bestpaths Changed: 0
      Time Spent: 0.000 secs

Last Notification Processing
      Received: 1d22h
      Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000005
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric       Notf      LastRIBEvent RefCount
10.0.101.201    [UR]            4294967295    0/0       1d22h (Reg)      0/3
94.0.0.2        [R][C][NL]            0        1/0       1d22h (Cri)     10/13


VRF: 905
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000006
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status          Metric       Notf      LastRIBEvent RefCount
10.0.101.201    [UR]            4294967295    0/0       1d22h (Reg)      0/3
95.0.0.2        [R][C][NL]            0        1/0       1d22h (Cri)     10/13


VRF: 906
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
```

```
    Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000007
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status        Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]          4294967295      0/0      1d22h (Reg)       0/3
96.0.0.2        [R][C][NL]             0      1/0      1d22h (Cri)      10/13


VRF: 907
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000008
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status        Metric        Notf      LastRIBEvent RefCount
10.0.101.201    [UR]          4294967295      0/0      1d22h (Reg)       0/3
97.0.0.2        [R][C][NL]             0      1/0      1d22h (Cri)      10/13


VRF: 908
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
```

```
   Received: 1d22h
   Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000009
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status         Metric         Notf      LastRIBEvent RefCount
10.0.101.201    [UR]           4294967295     0/0       1d22h (Reg)      0/3
98.0.0.2        [R][C][NL]              0      1/0       1d22h (Cri)     10/13


VRF: 909
========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe000000a
Nexthop Count: 1
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I    Invalid (Policy Match Failed)
Next Hop        Status         Metric         Notf      LastRIBEvent RefCount
99.0.0.2        [UR]           4294967295     0/0       1d22h (Reg)      0/3


VRF: yellow
===========

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 82y48w
  Time Spent: 0.000 secs

IPv4 Unicast is active
```

```
Gateway Address Family: IPv4 Unicast
Table ID: 0xe000000e
Nexthop Count: 0
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1
```

This table describes the significant fields shown in the display.

Table 24: show bgp vrf all nexthops Field Descriptions

Field	Description
VRF	Name of the VRF.
Total Nexthop Processing Time Spent	Time spent processing trigger delays for critical and noncritical events for the VRF or address family. The time is specified in seconds.
Maximum Nexthop Processing	Time that has passed since the nexthop notification was received that resulted in spending the maximum amount of processing time for all notifications.
Last Notification Processing	Time that has passed since the last nexthop notification was received.
IPv4 Unicast is active.	VRF specified output that indicates the IPv4 unicast address family is active within the VRF.
Nexthop Count	Number of next hops for the VRF or address family.
Critical Trigger Delay	Configured critical trigger delay.
Non-critical Trigger Delay	Configured noncritical trigger delay.
Total Critical Notifications Received	Number of critical notifications received.
Total Non-critical Notifications Received	Number of noncritical notifications received.
Bestpaths Deleted After Last Walk	Number of best paths deleted due to the last notification.
Bestpaths Changed After Last Walk	Number of best paths modified due to the last notification.
Next Hop	IP address of the next hop.
Status	Status of the next hop.
Metric	IGP metric of the next hop.

Field	Description
Notf	Number of critical and noncritical notifications received.
LastRIBEvent	When the last notification was received from the RIB.
RefCount	The number of neighbors or prefixes that refer to the next hop in address family/all format.
Address Family	Name of the address family.

Related Commands

Command	Description
bgp redistribute-internal, page 88	Specifies the delay for triggering BGP next-hop calculations.

show bgp nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information, use the **show bgp nsr** command in EXEC mode.

show bgp [**ipv4** {**unicast**| **multicast**| **labeled-unicast**| **all**| **tunnel**| **mdt**}| **ipv6** {**unicast**| **multicast**| **all**| **labeled-unicast**}| **all** {**unicast**| **multicast**| **all**| **labeled-unicast**| **mdt**| **tunnel**}| **vpnv4 unicast**| **vrf** {*vrf-name*| **all**} [**ipv4** {**unicast**| **labeled-unicast**}| **ipv6 unicast**]| **vpvn6 unicast**] **nsr** [**standby**]

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
ipv6	(Optional) Specifies IP Version 6 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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
standby	Displays information about the standby card.

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

Task ID

Task ID	Operations
bgp	read

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

```
RP/0/0/CPU0:router# show bgp nsr

Fri Jan 30 10:18:48.171 PST PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
Default Cluster ID: 10.1.0.1
Active Cluster IDs:  10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGPs notified
RIB has converged: version 1

========== Post Failover Summary for Active instance ==========
```

```
      Node            Process        Read     Write    Inbound

      node0_0_CPU0    Speaker        146.75    18.90     3.46

        Entered mode  Standby Ready                : Jan 30 10:00:39
        Entered mode  TCP NSR Setup                : Jan 30 10:00:39
        Entered mode  TCP NSR Setup Done           : Jan 30 10:00:39
        Entered mode  TCP Initial Sync             : Jan 30 10:00:39
        Entered mode  TCP Initial Sync Done        : Jan 30 10:00:44
        Entered mode  FPBSN processing done        : Jan 30 10:00:44
        Entered mode  Update processing done       : Jan 30 10:00:44
        Entered mode  BGP Initial Sync             : Jan 30 10:00:44
        Entered mode  BGP Initial Sync done        : Jan 30 10:00:44
        Entered mode  NSR Ready                    : Jan 30 10:00:44

      Current BGP NSR state - NSR Ready achieved at: Jan 30 10:00:44
      NSR State READY notified to Redcon at: Jan 30 10:16:58

      NSR Post Failover Summary:


      QAD Statistics:

        Messages Sent       : 512          ACKs Received       : 512
        Messages Received   : 8            ACKs Sent           : 8
        Send Failures       : 1            Send ACK Failures   : 0
        Suspends            : 1            Resumes             : 1
        Messages Processed  : 8            Out of sequence drops: 0

      Postit Summary:

        Total pending postit messages:  0
        Neighbors with pending postits: 0



      Conv  Bestpath   TunnelUpd  Import     RIBUpd     Label      ReadWrite  LastUpd
      Process: Speaker

      Yes   120        ---        ---        120        120        120        87531

      Rib Trigger: enabled
      Last RIB down event  Jan 29 09:50:03.069 received
      Last RIB convergence Jan 29 09:50:03.069 last ack received.

      Address Family IPv4 Unicast converged in 87531 seconds
```

The following example shows sample output from the **show bgp nsr** command with the **standby** keyword:

```
      RP/0/0/CPU0:router# show bgp nsr standby

      Fri Jan 30 10:18:55.654 PST PDT

      BGP Process Information:
      BGP is operating in STANDALONE mode
      Autonomous System: 100
      Router ID: 10.1.0.1 (manually configured)
      Default Cluster ID: 10.1.0.1
      Active Cluster IDs:  10.1.0.1
      Fast external fallover enabled
      Neighbor logging is not enabled
      Enforce first AS enabled
      AS Path ignore is enabled
      AS Path multipath-relax is enabled
      Default local preference: 100
      Default keepalive: 60
      Graceful restart enabled
      Restart time: 180
      Stale path timeout time: 360
      RIB purge timeout time: 600
      Non-stop routing is enabled
```

```
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGPs notified
RIB has converged: version 1

========== Post Failover Summary for Standby instance ==========

Node                Process            Read      Write     Inbound


node0_1_CPU0        Speaker            1.68      0.00      1.42

Entered mode  Standby Ready             : Jan 30 10:00:39
Entered mode  TCP Replication           : Jan 30 10:00:39
Entered mode  TCP Init Sync Done        : Jan 30 10:00:44
Entered mode  NSR Ready                 : Jan 30 10:00:44

QAD Statistics:

  Messages Sent      : 9        ACKs Received     : 9
  Messages Received  : 512      ACKs Sent         : 512
  Send Failures      : 0        Send ACK Failures  : 0
  Suspends           : 0        Resumes           : 0
  Messages Processed : 512      Standby init drops : 0       Out of sequence
drops: 0

Postit Summary:

  Total pending postit messages:  0
  Neighbors with pending postits: 0


Conv  Bestpath  TunnelUpd  Import    RIBUpd     Label     ReadWrite  LastUpd
Process: Speaker

Yes   1233338444 ---        ---       1233338444 1233338444 1233338444 ---

Rib Trigger: enabled
Last RIB down event  Jan 29 09:50:17.308 received
Last RIB convergence Jan 29 09:50:17.308 last ack received.
```

Related Commands

Command	Description
nsr (BGP), page 197	Activates Border Gateway Protocol (BGP) nonstop routing (NSR).

show bgp paths

To display all the Border Gateway Protocol (BGP) paths in the database, use the **show bgp paths** command in EXEC mode.

show bgp paths [detail] [debug] [**regexp** *regular-expression*]

Syntax Description

detail	(Optional) Displays detailed attribute information.
debug	(Optional) Displays attribute process ID, hash bucket, and hash chain ID attribute information.
regexp *regular-expression*	(Optional) Specifies an autonomous system path that matches the regular expression.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The **regexp** keyword was added.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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. See the *Cisco IOS XR Getting Started Guide for the Cisco XR 12000 Series Router* for information on regular expressions.

Use the **detail** keyword to display detailed information on the attributes stored with the AS path.

Task ID

Task ID	Operations
bgp	read

Examples

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

```
RP/0/0/CPU0:router# show bgp paths detail

Proc  Attributes                   Refcount    Metric Path
Spk 0 ORG AS LOCAL                      7          0 i
Spk 0 ORG AS LOCAL COMM EXTCOMM         3          0 21 i
Spk 0 MET ORG AS                        3         55 2 i
Spk 0 ORG AS                            3          0 2 10 11 i
Spk 0 ORG AS COMM                       3          0 2 10 11 i
Spk 0 MET ORG AS ATOM                   3          2 2 3 4 ?
Spk 0 MET ORG AS                        3          1 2 3 4 e
Spk 0 MET ORG AS                        3          0 2 3 4 i
```

This table describes the significant fields shown in the display.

Table 25: 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. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.

Field	Description
	• 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.
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 EXEC mode.

show bgp policy

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.	
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.	
ipv6	(Optional) Specifies IP Version 6 address prefixes.	
all	(Optional) For address family, specifies prefixes for all address families.	
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.	
vrf-name	(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.	
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.	
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.	
neighbor	(Optional) Previews advertisements for a single neighbor.	
ip-address	(Optional) IP address of a single neighbor.	
sent-advertisements	(Optional) Displays the routes that have been advertised to neighbors. If a route has not yet been advertised to the neighbor, it is not shown.	

route-policy	(Optional) Displays advertisements for an output route policy.
route-policy-name	(Optional) Name of the route policy.
standby	(Optional) Displays information about the standby card.
summary	(Optional) Displays a summary of the BGP advertisements.

Command Default

Advertisements for all neighbors are displayed if the **neighbor** *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

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **unsuppress-map** *map* keyword and argument were removed and the **route-policy** *route-policy-name* keyword and argument were added.
Release 3.3.0	The following keywords and arguments were added:
	• **vrf** { *vrf-name* | **all** }
	• [**ipv4** { **unicast** | **labeled-unicast** }]
	• [**vpnv4 unicast**] [**rd** *rd-address*]
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.
	The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp policy** command to display routes that would be advertised to neighbors under a proposed policy. Unlike in the **show bgp advertised** command, the information displayed reflects any modifications made to the routes when executing the specified policy.

Use the **neighbor** keyword to limit the output to routes advertised to a particular neighbor. Use the **sent-advertisements** keyword to change the output in two ways:

- If a policy is not specified explicitly, any policy configured on the neighbor (using the **route-policy (BGP)** command) is executed before displaying the routes.

- Only routes that have already been advertised to the neighbor (and not withdrawn) are displayed. Routes that have not yet been advertised are not displayed.

Use the **summary** keyword to display abbreviated output.

Task ID

Task ID	Operations
bgp	read

Examples The following is sample output from the **show bgp policy** command with the **summary** keyword in EXEC mode:

```
RP/0/0/CPU0:router# show bgp policy summary

Network          Next Hop        From            Advertised to
172.16.1.0/24    10.0.101.1      10.0.101.1      10.0.101.2
                                                 10.0.101.3

172.17.0.0/16    0.0.0.0         Local           10.0.101.1
                                                 10.0.101.2
                                                 10.0.101.3
```

This table describes the significant fields shown in the display.

Table 26: show bgp policy summary Field Descriptions

Field	Description
Network	IP prefix and prefix length for a network.

Field	Description
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
From	IP address of the peer that advertised this route.
Local	Indicates the route originated on the local system.
Local Aggregate	Indicates the route is an aggregate created on the local system.
Advertised to	Indicates the neighbors to which this route was advertised.

The following is sample output from the **show bgp policy** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp policy

11.0.0.0/24 is advertised to 10.4.101.1
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath: 1

11.0.0.0/24 is advertised to 10.4.101.2
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
    aspath:

11.0.0.0/24 is advertised to 10.4.101.3
 Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
    origin: IGP metric: 0
    aspath:
  Attributes after outbound policy was applied:
    next hop: 10.4.0.1
    MET ORG AS
    origin: IGP metric: 0
```

```
          aspath:

12.0.0.0/24 is advertised to 10.4.101.2
 Path info:
    neighbor: 10.4.101.1      neighbor router id: 10.4.101.1
    valid external best
  Attributes after inbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath: 2 3 4
  Attributes after outbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath:2 3 4

12.0.0.0/24 is advertised to 10.4.101.3
 Path info:
    neighbor: 10.4.101.1      neighbor router id: 10.4.101.1
    valid external best
  Attributes after inbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath: 2 3 4
  Attributes after outbound policy was applied:
    next hop: 10.4.101.1
    ORG AS
    origin: IGP neighbor as: 2
    aspath:2 3 4
```

This table describes the significant fields shown in the display.

Table 27: show bgp policy Field Descriptions

Field	Description
Is advertised to	IP address of the peer to which this route is advertised. If the route is advertised to multiple peers, information is shown separately for each peer.
neighbor	IP address of the peer that advertised this route, or one of the following: Local—Route originated on the local system. Local Aggregate—Route is an aggregate created on the local system.
neighbor router id	BGP identifier for the peer, or the local system if the route originated on the local system.
Not advertised to any peer	Indicates the no-advertise well-known community is associated with this route. Routes with this community are not advertised to any BGP peers.
Not advertised to any EBGP peer	Indicates the no-export well-known community is associated with this route. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.

Field	Description
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.
Attributes after inbound policy was applied	Displays attributes associated with the received route, after any inbound policy has been applied.

Field	Description
	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; in 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 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	Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities:

Field	Description
	Local-AS—Community with value 4294901812. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.
	no-advertise—Community with value 4294901813. Routes with this community value are not advertised to any BGP peers.
	no-export—Community with value 4294901814. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.
Extended community	Extended community attributes associated with the path. For known extended community types, the following codes may be displayed:
	RT—Route target community
	SoO—Site of Origin community
	LB—Link Bandwidth community
Originator	Router ID of the originating router when route reflection is used.
Cluster lists	Router ID or cluster ID of all route reflectors through which the route has passed.
Attributes after outbound policy was applied	Displays attributes associated with the received route, after any outbound policy has been applied.
	AGG—Aggregator attribute is present.
	AS—AS path attribute is present.
	ATOM—Atomic aggregate attribute is present.
	COMM—Communities attribute is present.
	EXTCOMM—Extended communities attribute is present.
	LOCAL—Local preference attribute is present.
	MET—Multi Exit Discriminator (MED) attribute is present.
	next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
	ORG—Origin attribute is present.

Related Commands

Command	Description
route-policy (BGP), page 225	Applies an inbound or outbound routing policy to a neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp advertised, page 256	Displays routes advertised to neighbors.
show bgp neighbors, page 310	Displays information about the TCP and BGP connections to neighbors.
show bgp route-policy, page 385	Displays BGP information about networks that match an outbound route policy.

show bgp process

To display Border Gateway Protocol (BGP) process information, use the **show bgp process** command in EXEC mode.

show bgp [**ipv4**| {**unicast**| **multicast**| **labeled-unicast**| **all**| **tunnel**| **mdt**}| **ipv6**| {**unicast**| **multicast**| **all**| **labeled-unicast**}| **all**| {**unicast**| **multicast**| **all**| **labeled-unicast**| **mdt**| **tunnel**}| **vpnv4 unicast**| **vpvn6 unicast**] **process** [**performance-statistics**] [**detail**] [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4.
unicast	(Optional) Specifies the unicast subaddress family.
multicast	(Optional) Specifies the multicast subaddress family.
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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6.
all	(Optional) For address family, specifies prefixes for all address families.
vpnv4 unicast	(Optional) Specifies VPNv4 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
performance- statistics	(Optional) Displays performance statistics relative to the work done by the specified process.
detail	(Optional) Specifies detailed process information.
standby	(Optional) Displays information about the standby card.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.4.0	The **vpnv4 unicast** and the **labeled-unicast** keywords were supported.
Release 3.5.0	The **vpnv6 unicast** keywords were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.
	The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.
Release 3.9.0	Asplain format for 4-byte autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.
Release 4.0	Output was modified to include information from BGP additional paths send and receive capability configurations.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Use the **show bgp process** command to display status and summary information for the Border Gateway Protocol (BGP) process. The output shows various global and address family-specific BGP configurations. A summary of the number of neighbors, update messages, and notification messages sent and received by the process is also displayed.

Use the **detail** keyword to display detailed process information. The detailed process information shows the memory used by each of various internal structure types.

Use the **performance-statistics** keyword to display a summary or detail of work done by the BGP processes. The summary display shows the real time spent performing certain operations and the time stamps for state transitions during initial convergence.

Task ID

Task ID	Operations
bgp	read

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

```
RP/0/0/CPU0:router# show bgp process

BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Scan interval: 60
Main Table Version: 150
IGP notification: IGPs notified

Node         Process    Nbrs Estab Rst Upd-Rcvd Upd-Sent Nfn-Rcvd Nfn-Sent
node0_0_CPU0 Speaker       3     2   1       20       10        0        0
```

This table describes the significant fields shown in the display.

Table 28: show bgp process Field Descriptions

Field	Description
BGP is operating in	Indicates BGP is operating in standalone mode. This is the only supported mode.
Autonomous System	Autonomous system number for the local system. • 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.
Router ID	BGP identifier assigned to the local system. If this is explicitly configured using the **bgp router-id** command, "manually configured" is displayed. If the router ID is not manually configured, it is determined

Field	Description
	from a global router ID. If no global ID is available, the router ID is shown as 0.0.0.0.
Confederation ID	Confederation identifier for the local system.
Cluster ID	Cluster identifier for the local system. If this is manually configured using the **bgp cluster-id** command, "manually configured" is displayed.
Default metric	Default metric. This is controlled by the **default-metric** command.
Fast external fallover enabled	Indicates whether fast external fallover is enabled. This is controlled by the **bgp fast-external-fallover disable** command.
Neighbor logging enabled	Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the **bgp log neighbor changes disable** command.
Enforce first AS enabled	Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled. This is controlled by the **bgp enforce-first-as disable** command.
iBGP to IGP redistribution	Indicates internal redistribution is enabled using the **bgp redistribution-internal** command.
Treating missing MED as worst	Indicates missing Multi Exit Discriminator (MED) metric values are treated as worst in the route selection algorithm. This is controlled by the **bgp bestpath med missing-as-worst** command.
Always compare MED is enabled	Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the **bgp bestpath med always** command.
AS Path ignore is enabled	Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the **bgp bestpath as-path ignore** command.
Comparing MED from confederation peers	Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the **bgp bestpath med confed** command.

Field	Description
Comparing router ID for eBGP paths	Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the **bgp bestpath compare-routerid** command.
Default local preference	Default local preference value used for BGP routes. This is controlled by the **bgp default local-preference** command.
Default keepalive	Default keepalive interval. This is controlled by the **timers bgp** command.
Graceful restart enabled	Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: • **bgp graceful-restart** • **bgp graceful-restart purge-time** • **bgp graceful-restart stalepath-time** • **bgp graceful-restart restart-time** • **bgp graceful-restart graceful-reset**
Update delay	Maximum time that a BGP process stays in read-only mode.
Generic scan interval	Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the **bgp scan-time** command.
Dampening	Indicates whether dampening is enabled for the specified address family. This is controlled by the **dampening** command.
Client reflection	Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the **bgp client-to-client reflection disable** command.
Scan interval	Interval (in seconds) between BGP scans for the given address family. This is controlled by the **bgp scan-time** command in address family configuration mode.
Main Table Version	Last version of the BGP database that was installed into the main routing table.

Cisco IOS XR Routing Command Reference for the Cisco XR 12000 Series Router, Release 4.1

Field	Description
IGP notification	Indicates whether Interior Gateway Protocols (IGP) have been notified of BGP convergence for the specified address family.
Node	Node on which the process is executing.
Process	Type of BGP process.
Speaker	Speaker process. A speaker process is responsible for receiving, processing, and sending BGP messages to configured neighbors.
Nbrs	Number of neighbors for which the process is responsible.
Estab	Number of neighbors that have connections in the established state for this process.
Rst	Number of times this process was restarted.
Upd-Rcvd	Number of update messages received by the process.
Upd-Sent	Number of update messages sent by the process.
Nfn-Rcvd	Number of notification messages received by the process.
Nfn-Sent	Number of notification messages sent by the process.

The following is sample output from the **show bgp process** command with the **detail** keyword:

```
RP/0/0/CPU0:router# show bgp all all process detail

BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60

BGP Speaker process: 0, location node0_0_0
Neighbors: 3, established: 2

                        Sent            Received
Updates:                3               15
Notifications:          0               0

                        Number          Memory Used
Attributes:             12              1104
AS Paths:               10              400
```

```
Communities:              2            1080
Extended communities:     1            40
Route Reflector Entries:  0            0
Route-map Cache Entries:  0            0
Filter-list Cache Entries: 0           0
Next Hop Cache Entries:    2           80
Update messages queued:    0

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Main Table Version: 12
IGP notification: IGPs notified

State: normal mode.
BGP Table Version: 12
Network Entries: 15, Soft Reconfig Entries: 0
Dampened Paths: 0, History Paths: 9

                          Allocated    Freed
Prefixes:                 15           0
Paths:                    19           0

                          Number       Memory Used
Prefixes:                 15           1230
Paths:                    19           760
```

This table describes the significant fields shown in the display.

Table 29: show bgp process detail Field Descriptions

Field	Description
BGP is operating in	Indicates whether BGP is operating in standalone mode.
Autonomous System	Autonomous system number for the local system.
Router ID	BGP identifier assigned to the local system. If this is explicitly configured using the **bgp router-id** command, "manually configured" is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0.
Confederation ID	Confederation identifier for the local system.
Cluster ID	Cluster identifier for the local system. If this is manually configured using the **bgp cluster-id** command, "manually configured" is displayed.
Default metric	Default metric.
Fast external fallover enabled	Indicates whether fast external fallover is enabled.
Neighbor logging enabled	Indicates whether logging of peer connection up and down transitions is enabled.

Field	Description
Enforce first AS enabled	Indicates that strict checking of the first autonomous system (AS) number in paths received from external BGP peers is enabled.
iBGP to IGP redistribution	Indicates internal redistribution is enabled using the **bgp redistribution-internal** command.
Treating missing MED as worst	Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled by the **bgp bestpath med missing-as-worst** command.
Always compare MED is enabled	Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the **bgp bestpath med always** command.
AS Path ignore is enabled	Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the **bgp bestpath as-path ignore** command.
Comparing MED from confederation peers	Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the **bgp bestpath med confed** command.
Comparing router ID for eBGP paths	Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the **bgp bestpath compare-routerid** command.
Default local preference	Default local preference value used for BGP routes.
Default keepalive	Default keepalive interval. This is controlled by the **timers bgp** command.
Graceful restart enabled	Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: • **bgp graceful-restart** • **bgp graceful-restart purge-time** • **bgp graceful-restart stalepath-time** • **bgp graceful-restart restart-time** • **bgp graceful-restart graceful-reset**

Field	Description
Update delay	Maximum time that a BGP process stays in read-only mode.
Generic scan interval	Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the **bgp scan-time** command.
BGP Speaker Process	Speaker process responsible for receiving, processing and sending BGP messages.
Node	Node on which the specified process is executing.
Neighbors	Number of neighbors for which the specified process is responsible.
established	Number of neighbors that have connections in the established state for the specified process.
Updates	Number of update messages sent and received by the specified process.
Notifications	Number of notification messages sent and received by the specified process.
Attributes	Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information.
AS Paths	Number of unique autonomous system paths stored in the specified process and the amount of memory used by the AS path information.
Communities	Number of unique sets of community information stored in the specified process and the amount of memory used by them.
Extended communities	Number of unique sets of extended community information stored in the specified process and the amount of memory used by them.
Route Reflector Entries	Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them.
Nexthop Entries	Number of entries and memory usage for cached next-hop information.

Field	Description
Update messages queued	Total number of update messages queued to be sent across all neighbors for which the specified process is responsible.
Address family	Specified address family.
Dampening	Indicates whether dampening is enabled for the specified address family.
Client reflection	Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the **bgp client-to-client reflection disable** command.
Scan interval	Interval (in seconds) between BGP scans for the given address family. This is controlled by the **bgp scan-time** command.
Main Table Version	Last version of the local BGP database for the specified address family that was injected into the main routing table.
IGP notification	Indicates whether IGPs have been notified of BGP convergence for the specified address family.
RIB has converged	Indicates whether the main routing table version has converged and the version at which it converged.
State	BGP system state for the specified address family and process. This may be one of the following:
	read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers.
	best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode.
	import mode—Routes are imported from one VRF to another VRF once the best paths are calculated. This mode is supported in VPNv4 unicast address family mode.
	RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB.
	label allocation mode: Labels are allocated for the received prefixes based on the requirement.

Field	Description
	normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.
BGP Table Version	Last version used in the BGP database for received routes.
Attribute download	Indicates whether the RIB attribute download is enabled.
Network Entries	Number of sets of prefix information held in the specified BGP process for the specified address family.
Soft Reconfig Entries	Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration.
Dampened Paths	Number of routes that are suppressed due to dampening for the specified address family.
History Paths	Number of routes that are currently withdrawn, but are being maintained to preserve dampening information.
Prefixes (Allocated/Freed)	Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process.
Paths (Allocated/Freed)	Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process.
Prefixes (Number/Memory Used)	Number of sets of prefix information currently allocated for the specified address family, and the amount of memory used by them.
Paths (Number/Memory Used)	Number of sets of route information currently allocated for the specified address family, and the amount of memory used by them.

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

```
RP/0/0/CPU0:router# show bgp process performance-statistics detail

BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2
```

```
                         Sent           Received
Updates:                 20             20
Notifications:           0              0

                         Number         Memory Used
Attributes:              2              184
AS Paths:                2              48
Communities:             0              0
Extended communities:    0              0
Route Reflector Entries: 0              0
Route-map Cache Entries: 0              0
Filter-list Cache Entries: 0            0
Next Hop Cache Entries:   2             80
Update messages queued:  0

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)


State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0



                         Allocated      Freed
Prefixes:                149            0
Paths:                   200            0

                         Number         Memory Used
Prefixes:                149            12516
Paths:                   200            8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s
```

This table describes the significant fields shown in the display.

Table 30: show bgp process performance-statistics Field Descriptions

Field	Description
BGP is operating in	Indicates whether BGP is operating in standalone mode.
Autonomous system	Autonomous system number for the local system.
Router ID	BGP identifier assigned to the local system. If this is explicitly configured using the **bgp router-id** command, "manually configured" is displayed. If the router ID is not manually configured, it is determined

Field	Description
	from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0.
Confederation ID	Confederation identifier for the local system.
Cluster ID	The cluster identifier for the local system. If this is manually configured using the **bgp cluster-id** command, "manually configured" is displayed.
Default metric	Default metric.
Fast external fallover enabled	Indicates whether fast external fallover is enabled.
Neighbor logging enabled	Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the **bgp log neighbor changes disable** command.
Enforce first AS enabled	Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled.
iBGP to IGP redistribution	Indicates internal redistribution is enabled using the **bgp redistribution-internal** command.
Treating missing MED as worst	Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled using the **bgp bestpath med missing-as-worst** command.
Always compare MED is enabled	Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This setting is controlled by the **bgp bestpath med always** command.
AS Path ignore is enabled	Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the **bgp bestpath as-path ignore** command.
Comparing MED from confederation peers	Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the **bgp bestpath med confed** command.
Comparing router ID for eBGP paths	Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the **bgp bestpath compare-routerid** command.

Field	Description
Default local preference	Default local preference value used for BGP routes.
Default keepalive	Default keepalive interval. This setting is controlled by the **timers bgp** command.
Graceful restart enabled	Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: **bgp graceful-restart**, **bgp graceful-restart purge-time**, **bgp graceful-restart stalepath-time**, **bgp graceful-restart restart-time**, and **bgp graceful-restart graceful-reset**.
Update delay	Maximum time that a BGP process stays in read-only mode.
Generic scan interval	Interval (in seconds) between BGP scans for address family-independent tasks. This setting is controlled by the **bgp scan-time** command in router configuration mode.
Address family	Specified address family.
Dampening	Indicates whether dampening is enabled for the specified address family.
Client reflection	Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the **bgp client-to-client reflection disable** command.
Scan interval	Interval (in seconds) between BGP scans for the given address family. This is controlled by the **bgp scan-time** command.
Main Table Version	Last version of the local BGP database for the specified address family that was injected into the main routing table.
IGP notification	Indicates whether IGPs have been notified of BGP convergence for the specified address family.
Node	Node on which the process is executing.
Process	BGP process.
Speaker	Speaker process. The speaker process is responsible for receiving, processing and sending BGP messages.

Field	Description
Read	Real time (in seconds) spent reading messages from peers by this process.
Write	Real time (in seconds) spent writing messages to peers by this process.
Inbound	The real time (in seconds) spent processing messages read from peers by this process.
Config	Real time (in seconds) spent processing configuration commands by this process.
Data	Real time (in seconds) spent providing operational data by this process.
Conv	Indicates whether the process has converged after the initial update.
Nbr Estab	Time stamp (in seconds) recording the time when the first neighbor became established.
Bestpath	Time stamp (in seconds) recording the time the best-path calculation mode was entered.
RIB Inst	Time stamp (in seconds) recording the time RIB update mode was entered.
Read/Write	Time stamp (in seconds) recording the time normal mode was entered.
Last Upd	Time stamp (in seconds) recording the time the last update was sent to a neighbor.
Address Family IPv4 Unicast converged in *n* seconds	Indicates that BGP has reached initial convergence for the IPv4 unicast address family. The time taken for convergence is shown.
Address Family IPv6 Multicast converged in *n* seconds	Indicates that BGP has reached initial convergence for the IPv6 multicast address family. The time taken for convergence is shown.

The following is sample output from the **show bgp process** command with the **performance-statistics** and **detail** keywords:

```
RP/0/0/CPU0:router# show bgp process performance-statistics detail

BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2
```

```
                              Sent            Received
Updates:                      20              20
Notifications:                0               0

                              Number          Memory Used
Attributes:                   2               184
AS Paths:                     2               48
Communities:                  0               0
Extended communities:         0               0
Route Reflector Entries:      0               0
Route-map Cache Entries:      0               0
Filter-list Cache Entries: 0                  0
Next Hop Cache Entries:        2               80
Update messages queued:       0

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)


State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0



                              Allocated       Freed
Prefixes:                     149             0
Paths:                        200             0

                              Number          Memory Used
Prefixes:                     149             12516
Paths:                        200             8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s
```

This table describes the significant fields shown in the display.

Table 31: show bgp process performance-statistics detail Field Descriptions

Field	Description
Process	The specified process.
Location	Node in which the specified process is executing.
Neighbors	Number of neighbors for which the specified process is responsible.
established	Number of neighbors that have connections in the established state for the specified process.

Field	Description
Updates	Number of update messages sent and received by the specified process.
Notifications	Number of notification messages sent and received by the specified process.
Attributes	Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information.
AS Paths	Number of unique autonomous system paths stored in the specified process, and the amount of memory used by the AS path information.
Communities	Number of unique sets of community information stored in the specified process and the amount of memory used by them.
Extended communities	Number of unique sets of extended community information stored in the specified process and the amount of memory used by them.
Route Reflector Entries	Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them.
Route-map Cache Entries	Number of entries and memory usage for cached results for applying a route map.
Filter-list Cache Entries	Number of entries and memory usage for cached results for applying an AS path filter list.
Next Hop Cache Entries	Number of entries and memory usage for cached next-hop information.
Update messages queued	Number of update messages queued to be sent across all neighbors for which the specified process is responsible.
Read	Indicates the number of messages read by the process, the total size of read messages, the number of read operations performed, and the real time spent by the process performing read operations.
Read throttled	Number of times that reading from TCP has been throttled due to a backlog of messages read but not processed.

Field	Description
inbound messages	Number of read messages that have been processed and the real time spent processing inbound messages.
Wrote	Amount of data that has been written by the process, the number of write operations performed, and the real time spent by the process performing write operations.
Processing write list	Number of messages written from write lists, the number of times the write list has been processed, and the real time spent processing the write list. **Note** Write lists typically contain only update messages.
Processing write queue	Number of messages written from write queues, number of times the write queue has been processed, and the real time spent processing the write queue.
Socket setup	Number of socket setup operations performed and the real time spent during socket setup operations.
Configuration	Number of configuration requests received by the process and the real time spent processing configuration requests.
Operational data	Number of requests for operational data (for **show** commands) received by the process and the real time spent processing operation data requests
State	BGP system state for the specified address family and process. This may be one of the following: read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers. best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode. import mode—Routes are imported from one VRF to another VRF once the best paths are calculated. This mode is supported in VPNv4 unicast address family mode. RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB. label allocation mode: Labels are allocated for the received prefixes based on the requirement.

Field	Description
	normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.
BGP Table Version	Last version used in the BGP database for received routes.
Network Entries	Number of sets of prefix information held in the specified BGP process for the specified address family.
Soft Reconfig Entries	Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration.
Dampened Paths	Number of routes that are suppressed due to dampening for the specified address family.
History Paths	Number of routes that are currently withdrawn, but are being maintained to preserve dampening information.
Prefixes (Allocated/Freed)	Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process.
Paths (Allocated/Freed)	Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process.
Prefixes (Number/Memory Used)	Number of sets of prefix information currently allocated for the specified address family and amount of memory used by them.
Paths (Number/Memory Used)	Number of sets of route information currently allocated for the specified address family and amount of memory used by them.
Updates generated	Number of prefixes for which updates have been generated, the number of messages used to advertise the updates, the number of update generation runs performed, and the real time spent generating updates for the specified address family.
Scanner	Number of times the scanner has run for the specified address family and real time spent in scanner processing.

Field	Description
RIB Update	Number of global routing information base update runs performed for the specified address family, number of prefixes installed, withdrawn, or modified in the global RIB during these runs, and real time spent performing these runs.
Process has converged	Indicates whether the process has reached initial convergence for the specified address family.
First neighbor established	Time stamp (in seconds) recording the time the first neighbor in the process was established.
Entered DO_BESTPATH mode	Time stamp (in seconds) recording the time best-path calculation mode was entered.
Entered DO_RIBUPD mode	Time stamp (in seconds) recording the time RIB update mode was entered.
Entered Normal mode	Time stamp (in seconds) recording the time normal mode was entered.
Last UPDATE sent	Time stamp (in seconds) recording the time the last update was sent to a neighbor.

The following is sample output from the **show bgp vpnv4 unicast process performance-statistics detail** command:

```
RP/0/0/CPU0:router# show bgp vpnv4 unicast process performance-statistics detail
BGP Speaker process: 0, Node: node0_8_CPU0 Restart count: 1
                        Total           Nbrs Estab/Cfg
Default VRFs:           1               4/12
Non-Default VRFs:       1009            1082/1337

                        Sent            Received
Updates:                362259          5688505
Notifications:          14              0

                        Number          Memory Used
Attributes:             14896           2979200
AS Paths:               17              1100
Communities:            3               120
Extended communities:   1849            124440
Route Reflector Entries: 417            25020
Nexthop Entries:        2941            539572
Update messages queued: 0

                        Alloc           Free
Pool 210:               28955629        28955628
Pool 310:               363103          363103
Pool 600:               4931162         4931162
Pool 1100:              104693          104693
Pool 4300:              799374          799374

Read 34755745 messages (3542094326 bytes) in 30528983 calls (time spent: 6427.769 secs)
Read partly throttled 1506 times
   Read 14 times after crossing lower threshold Processed 5836892 inbound update messages
 (time spent: 6229.512 secs) Wrote 825719955 bytes in 29272669 calls (time spent: 2318.472
 secs) Processing sub-group: wrote 861402 messages in 1113810 calls (time spent: 145.446
```

```
secs) Processing write queue: wrote 6288 messages in 20498 calls (time spent: 0.039 secs)
Socket setup (LPTS): 0 calls (time spent: 0.000 secs) event_file_attach calls: Input 8769,
 Output 2810, Input-output 0
Configuration: 989 requests (time spent: 0.046 secs) Operational data: 92396 requests (time
 spent: 98.864 secs) Current Clock Time: not set Update Generation master timer:
    id: 0, time left: 0.0 sec, last processed: not set
    expiry time of parent node: not set
IO master timer:
    id: 0, time left: 0.0 sec, last processed: not set
    expiry time of parent node: not set


Address Family: VPNv4 Unicast
State: Normal mode.
BGP Table Version: 23211188
Attribute download: Disabled
Soft Reconfig Entries: 0

                    Last 8 Triggers      Ver        Tbl Ver

Label Thread        Jun 18 05:31:39.120  23211188   23211188
                    Jun 18 05:31:35.274  23211188   23211188
                    Jun 18 05:31:34.340  23211187   23211188
                    Jun 18 05:31:34.189  23211186   23211187
                    Jun 18 05:31:29.120  23211186   23211186
                    Jun 18 05:31:28.861  23211186   23211186
                    Jun 18 05:31:19.640  23211186   23211186
                    Jun 18 05:31:19.272  23211186   23211186
                    Total triggers: 639526


Import Thread       Jun 18 05:31:39.120  23211188   23211188
                    Jun 18 05:31:35.274  23211188   23211188
                    Jun 18 05:31:34.340  23211187   23211188
                    Jun 18 05:31:34.189  23211186   23211187
                    Jun 18 05:31:29.120  23211186   23211186
                    Jun 18 05:31:28.861  23211186   23211186
                    Jun 18 05:31:19.640  23211186   23211186
                    Jun 18 05:31:19.272  23211186   23211186
                    Total triggers: 689177


RIB Thread          Jun 18 05:31:39.146  23211188   23211188
                    Jun 18 05:31:35.299  23211188   23211188
                    Jun 18 05:31:34.525  23211187   23211188
                    Jun 18 05:31:34.494  23211186   23211188
                    Jun 18 05:31:34.340  23211186   23211188
                    Jun 18 05:31:34.255  23211186   23211188
                    Jun 18 05:31:29.146  23211186   23211186
                    Jun 18 05:31:28.886  23211186   23211186
                    Total triggers: 668084


Update Thread       Jun 18 05:31:39.171  ---        23211188
                    Jun 18 05:31:35.324  ---        23211188
                    Jun 18 05:31:34.558  ---        23211188
                    Jun 18 05:31:34.521  ---        23211188
                    Jun 18 05:31:34.327  ---        23211188
                    Jun 18 05:31:29.170  ---        23211186
                    Jun 18 05:31:28.910  ---        23211186
                    Jun 18 05:31:19.690  ---        23211186
                    Total triggers: 660143

                    Allocated       Freed
Remote Prefixes:    3150972         2885064
Remote Paths:       7639074         7118286

Local Prefixes:     3760870         3425614
Local Paths:        7892100         7595657

                    Number          Mem Used
Remote Prefixes:    265908          29781696
Remote Paths:       520788          24997824
Remote RDs:         12424           2832672

Local Prefixes:     335256          37548672
```

```
Local Paths:        296443          14229264
Local RDs:          1009            230052

Total Prefixes:     601164          67330368
Total Paths:        817231          39227088
Imported Paths:     265675          12752400
Total RDs:          13433           3062724
Same RDs:           0               0


Update Groups: 3  Subgroups: 2
Updates generated: 1438448 prefixes in 67375 messages from 181564 calls (time spent: 6779.576
 secs)
Scanner: 0 scanner runs (time spent: 0.000 secs) RIB update: 0 rib update runs, 0 prefixes
 installed, 0 modified,
         0 prefixes removed (time spent: 0.000 secs) RIB table update: 0 table deletes,
 0 table invalid, 3526736604 table skip,
         0 no local label, 0 rib retries Process has not converged for VPNv4 Unicast.

First neighbor established:  Jun 11 08:32:10
Entered DO_BESTPATH mode:    Jun 11 08:52:10
Entered DO_IMPORT mode:      Jun 11 08:52:12
Entered DO_LABEL_ALLOC mode: Jun 11 08:52:16
Entered DO_RIBUPD mode:      Jun 11 08:52:19
Entered Normal mode:         Jun 11 08:52:23
Latest UPDATE sent:          Jun 18 05:31:34
```

The following is sample output from show bgp process detail command with information on additional paths send and receive information:

```
BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASDOT
Autonomous System: 100
Router ID: 22.22.22.22 (manually configured)
Default Cluster ID: 2.2.2.2 (manually configured)
Active Cluster IDs:  2.2.2.2
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 120
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

……
……
                    Allocated     Freed
Prefixes:           12            0
Paths:              60            0
Path-elems:         12            0

                    Number        Mem Used
Prefixes:           12            1200
Paths:              60            3120
Path-elems:         12            624
```

Related Commands

Command	Description
bgp bestpath as-path ignore, page 44	Sets the autonomous system path length to ignore when calculating preferred paths.

BGP Commands on Cisco IOS XR Software

show bgp process

Command	Description
bgp bestpath compare-routerid, page 46	Compare identical routes received from external BGP (eBGP) peers during the best-path selection process and select the route with the lowest router ID.
bgp bestpath med always, page 50	Compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
bgp bestpath med missing-as-worst, page 54	Assume paths with no MED attribute have the most undesirable MED value possible when performing path selection.
bgp cluster-id, page 58	Enables reflection of routes between route reflector clients using a BGP route reflector.
bgp cluster-id, page 58	Configure the cluster ID if the BGP cluster has more than one route reflector.
bgp default local-preference, page 67	Sets the default local preference value.
bgp redistribute-internal, page 88	Allows the redistribution of iBGP routes into an IGP such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF).
bgp router-id, page 90	Configures a fixed router ID for a BGP-speaking router.
default-metric (BGP), page 134	Sets default metric values for the BGP.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
bgp scan-time, page 92	Configures scanning intervals.
timers bgp, page 448	Sets default BGP timers.

Cisco IOS XR Routing Command Reference for the Cisco XR 12000 Series Router, Release 4.1

OL-24679-01

379

show bgp regexp

To display routes matching the autonomous system path regular expression, use the **show bgp regexp** command in EXEC mode.

show bgp regexp *regular-expression*

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
regular-expression	Regular expression to match the BGP autonomous system paths.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The following keywords and argument were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }] • **vpnv4 unicast**
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** keywords were added. The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp regexp** command to display all routes in the specified BGP table whose autonomous system path is matched by the specified regular expression.

Note   If the regular expression contains spaces and parentheses, it must be specified and surrounded by quotation marks.

Task ID

Task ID	Operations
bgp	read

Examples

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

```
RP/0/0/CPU0:router# show bgp regexp "^3 "

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 64
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network      Next Hop          Metric LocPrf Weight Path
*>i172.20.17.121    10.0.101.2                    100      0 3 2000 3000 i
*>i10.0.0.0         10.0.101.2                    100      0 3 100 1000 i
*>i172.5.23.0/24    10.0.101.2                    100      0 3 4 60 4378 i
```

This table describes the significant fields shown in the display.

Table 32: show bgp regexp Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
Dampening enabled	Displayed if dampening has been enabled for the routes in this BGP routing table.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	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): 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. The second character may be (in order of precedence):

Field	Description
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is placed at the end of each line in the table. It can be one of the following values:
	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.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network.
Metric	Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric.
LocPrf	Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system.
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

Related Commands

Command	Description
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, page 244	Displays entries in the BGP routing table.
show bgp route-policy, page 385	Displays BGP information about networks that match an outbound route policy.

show bgp route-policy

To display Border Gateway Protocol (BGP) information about networks that match an outbound route policy, use the **show bgp route-policy** command in EXEC mode.

show bgp route-policy *route-policy-name* [**standby**]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 address prefixes.
unicast	(Optional) Specifies unicast address prefixes.
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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
route-policy-name	Name of a route policy.
standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **count-only** keyword was added.
Release 3.3.0	The following keywords and arguments were added:
	• **vrf** { *vrf-name* | **all** }
	• [**ipv4** { **unicast** | **labeled-unicast** }]
	• [**vpnv4 unicast**] [**rd** *rd-address*]
Release 3.4.0	The **labeled-unicast** keyword was supported.
	The **count-only** keyword was removed.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.
	The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined.

If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

A route policy must be configured to use this command. When the **show bgp route-policy** command is entered, routes in the specified BGP table are compared with the specified route policy, and all routes passed by the route policy are displayed.

If a pass clause is encountered while the route policy is being applied to the route and the route policy processing completes without hitting a drop clause, the route is displayed. The route is not displayed if a drop clause is encountered, if the route policy processing completes without hitting a pass clause, or if the specified route policy does not exist.

The information displayed does not reflect modifications the policy might make to the route. To display such modifications, use the **show bgp policy** command.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp route-policy** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp route-policy p1

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 729
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*  10.13.0.0/16     192.168.40.24                        0 1878 704 701 200 ?
*  10.16.0.0/16     192.168.40.24                        0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 33: show bgp route-policy Field Descriptions

Field	Description
BGP router identifier	BGP identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.

Field	Description
Status codes	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): 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. The second character may be (in order of precedence): >—Path is the best path to use for that network. 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. The third character may be: i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** 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.
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.

Related Commands

Command	Description
aggregate-address, page 24	Configures an aggregate entry in a BGP routing table.
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor
route-policy	Configures a route policy.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp policy, page 348	Displays advertisements under a proposed policy.

show bgp session-group

To display information about the Border Gateway Protocol (BGP) configuration for session groups, use the **show bgp session-group** command in EXEC mode.

show bgp session-group *group-name* {**configuration [defaults] [nvgen]| inheritance| users**}

Syntax Description

group-name	Name of the session family group to display.
configuration	(Optional) Displays the effective configuration for the session group, including any inherited configuration.
defaults	(Optional) Displays all configuration, including default configuration.
nvgen	(Optional) Displays output in the form of the **show running-config** command. If the **defaults** keyword also is specified, the output is not suitable for cutting and pasting into a configuration session.
inheritance	(Optional) Displays the session groups from which this session group inherits configuration.
users	(Optional) Display the session groups, neighbor groups, and neighbors that inherit configuration from this session group.

Command Default No default behavior or value

Command Modes EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes 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 session-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a session group, including any configuration inherited from other session groups through application of the **use** command. The source for each configured command is also displayed.

Use the **defaults** keyword to display the value of all configuration, including default configuration. Use the **nvgen** keyword to display configuration in the form of the **show running-config** command output. Output in this form is suitable for cutting and pasting into a configuration session.

Use the **show bgp session-group** command with the *group-name* **inheritance** argument and keyword to display the session groups from which the specified session group inherits configuration.

Use the **show bgp session-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and session groups that inherit configuration from the specified session group.

Task ID

Task ID	Operations
bgp	read

Examples

For the example shown here, the following configuration is used:

```
session-group group3
  advertisement-interval 5
  dmzlink-bw
 !
 session-group group1
  use session-group group2
  update-source Loopback0
 !
 session-group group2
  use session-group group3
  ebgp-multihop 2
```

The following example shows the **show bgp session-group** command with the **configuration** keyword:

```
RP/0/0/CPU0:router# show bgp session-group group1 configuration

session-group group1
 advertisement-interval 5[s:group2 s:group3]
 ebgp-multihop 2          [s:group2]
 update-source Loopback0 []
 dmzlink-bandwidth              [s:group2 s:group3]
```

The source of each command is shown to the right of the command. For example, **update-source** is configured directly on session group group1. The **dmzlink-bandwidth** command is inherited from session group group2, which in turn inherits it from session group group3.

The following example shows the **show bgp session-group** command with the **users** keyword:

```
RP/0/0/CPU0:router# show bgp session-group group2 users

IPv4 Unicast:a:group1
```
The following example shows the **show bgp session-group** command with the **inheritance** keyword.

```
RP/0/0/CPU0:router# show bgp session-group group1 inheritance

Session:s:group2 s:group3
```

The command output shows that the session group group1 directly uses the group2 session group. The group2 session group uses the group3 session group.

This table describes the significant fields shown in the display.

Table 34: show bgp session-group Field Descriptions

Field	Description
[ ]	Configures the command directly on the specified session group.
s:	Indicates the name that follows is a session group.
a:	Indicates the name that follows is an address family group.
n:	Indicates the name that follows is a neighbor group.
[dflt]	Indicates the command is not explicitly configured or inherited, and the default value for the command is used. This field may be shown when the **defaults** keyword is specified.
<not set>	Indicates that the default is for the command to be disabled. This field may be shown when the **defaults** keyword is specified.

Related Commands

Command	Description
session-group, page 240	Configures a BGP session group.
show bgp neighbor-group, page 306	Displays information about the BGP configuration for neighbor groups.
show bgp neighbors, page 310	Displays information about BGP connections to neighbors.

show bgp sessions

To display brief information about BGP neighbors, use the **show bgp sessions** command in EXEC mode.

show bgp sessions [not-established] [not-nsr-ready]

Syntax Description

not-established	(Optional) Displays all the neighbors that are not in established state
not-nsr-ready	(Optional) Displays all the neighbors that are not nonstop routing (NSR) ready.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

The **show bgp sessions** command without a keyword provides brief information about all the BGP neighbors configured irrespective of the address family or VRF.

The **show bgp sessions** command with the **not-established** keyword shows BGP peers which are yet to establish their peering relationship.

The **show bgp session**s command with the and **not-nsr-ready** keyword shows BGP peers which are yet to reach the nsr ready state.

Task ID

Task ID	Operations
bgp	read

Examples   The following is sample output from the **show bgp sessions** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp sessions
Thu Jan 15 17:41:45.277 UTC

Neighbor      VRF            Spk   AS   InQ  OutQ  NBRState    NSRState
2.2.2.2       default         0    1    0     0   Active      None
10.0.101.1    default         0    1    0     0   Established  NSR Ready
```

```
10.0.101.2      default                 0     1    0     0  Established  NSR Ready
10.0.101.3      default                 0     1    0     0  Established  NSR Ready
10.0.101.4      default                 0     1    0     0  Established  NSR Ready
10.0.101.5      default                 0     1    0     0  Established  NSR Ready
10.0.101.6      default                 0     1    0     0  Established  NSR Ready
10.0.101.7      default                 0     1    0     0  Established  NSR Ready
10.0.101.8      default                 0     1    0     0  Established  NSR Ready
10.0.101.9      default                 0     1    0     0  Established  NSR Ready
10.11.12.2      default                 0   100    0     0  Established  NSR Ready
90.0.0.2        900                     0     2    0     0  Established  NSR Ready
9000::1001      900                     0     2    0     0  Established  NSR Ready
91.0.0.2        901                     0     2    0     0  Established  NSR Ready
9100::1001      901                     0     2    0     0  Established  NSR Ready
92.0.0.2        902                     0     2    0     0  Established  NSR Ready
9200::1001      902                     0     2    0     0  Established  NSR Ready
93.0.0.2        903                     0     2    0     0  Established  NSR Ready
9300::1001      903                     0     2    0     0  Established  NSR Ready
94.0.0.2        904                     0     2    0     0  Established  NSR Ready
9400::1001      904                     0     2    0     0  Established  NSR Ready
95.0.0.2        905                     0     2    0     0  Established  NSR Ready
9500::1001      905                     0     2    0     0  Established  NSR Ready
96.0.0.2        906                     0     2    0     0  Established  NSR Ready
9600::1001      906                     0     2    0     0  Established  NSR Ready
97.0.0.2        907                     0     2    0     0  Established  NSR Ready
9700::1001      907                     0     2    0     0  Established  NSR Ready
98.0.0.2        908                     0     2    0     0  Established  NSR Ready
9800::1001      908                     0     2    0     0  Established  NSR Ready
99.0.0.2        909                     0     2    0     0  Idle         None
9900::1001      909                     0     2    0     0  Idle         None
12.13.14.16     red                     0     2    0     0  Idle         None
20.0.101.1      red                     0     2    0     0  Active       None
1234:5678:9876::1111
                red                     0     3    0     0  Idle         None
2020::1002      red                     0     2    0     0  Established  NSR Ready
1.2.3.4         this-is-a-long-vrf-name
                                        0     5    0     0  Idle         None
1111:2222:3333:4444:5555::6789
                this-is-a-long-vrf-name
                                        0     7    0     0  Idle         None
```

The following is sample output from the **show bgp sessions** command with the **not-established** keyword:

```
RP/0/0/CPU0:router# show bgp sessions not-established
Fri Jan 30 11:30:42.720 PST PDT

Neighbor        VRF             Spk    AS   InQ  OutQ  NBRState     NSRState
10.0.101.5      default           0   100    0     0  Active       None
2.2.2.2         vrf1_1           0   302    0     0  Idle         None
2.101.1.2       vrf1_1           0   302    0     0  Idle         None
2.102.1.2       vrf1_1           0   302    0     0  Idle         None
2.103.1.2       vrf1_1           0   302    0     0  Idle         None
4.4.4.2         vrf1_1           0   304    0     0  Idle         None
2008:2:2:2::2   vrf1_1           0   302    0     0  Idle         None
11.16.1.2       vrf2_1           0   302    0     0  Idle         None
```

The following is sample output from the **show bgp sessions** command with the **not-nsr-ready** keyword:

```
RP/0/0/CPU0:router# show bgp sessions not-nsr-ready
Fri Jan 30 11:30:52.301 PST PDT

Neighbor        VRF             Spk    AS   InQ  OutQ  NBRState     NSRState
10.0.101.5      default           0   100    0     0  Active       None
2.2.2.2         vrf1_1           0   302    0     0  Idle         None
2.101.1.2       vrf1_1           0   302    0     0  Idle         None
2.102.1.2       vrf1_1           0   302    0     0  Idle         None
2.103.1.2       vrf1_1           0   302    0     0  Idle         None
4.4.4.2         vrf1_1           0   304    0     0  Idle         None
2008:2:2:2::2   vrf1_1           0   302    0     0  Idle         None
11.16.1.2       vrf2_1           0   302    0     0  Idle         None
```

This table describes the significant fields shown in the display.

Table 35: show bgp sessions Field Descriptions

Field	Description
Neighbor	Displays neighbor IP address.
VRF	Displays information about the VRF.
Spk	Speaker process that is responsible for the neighbor. Always 0.
AS	Autonomous system.
InQ	Number of messages from a neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to a neighbor.
NBRState	State of the Border Gateway Protocol (BGP) neighbor sessions.
NSRState	State of the Border Gateway Protocol (BGP) nonstop routing (NSR).

Related Commands

Command	Description
show bgp neighbors, page 310	Displays information about Border Gateway Protocol (BGP) connections to neighbors.

show bgp summary

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

show bgp [ipv4 {**unicast**| **multicast**| **labeled-unicast**| **all**| **tunnel**| **mdt**}| **ipv6** {**unicast**| **multicast**| **all**| **labeled-unicast**}| **all** {**unicast**| **multicast**| **all**| **labeled-unicast**| **mdt**| **tunnel**}| **vpnv4 unicast**| **vrf** {*vrf-name*| **all**} [**ipv4** {**unicast**| **labeled-unicast**}| **ipv6 unicast**]| **vpvn6 unicast**] **summary** [**standby**]

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
ipv6	(Optional) Specifies IP Version 6 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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The following keywords and argument were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }] • **vpnv4 unicast**
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** keywords were added. The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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. See *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Use the **show bgp summary** command to display a summary of the neighbors for which the specified address family and subaddress family are enabled. If the neighbor does not have the specified address family and subaddress family enabled, it is not included in the output of the **show** command. If the **all** keyword is specified for the address family or subaddress family, a summary for each combination of address family and subaddress family is displayed in turn.

The table versions shown in the output (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 < RecvTblVer—Some received routes have not yet been considered for installation in the global routing table.

- TblVer < SendTblVer—Some received routes have been installed in the global routing table but have not yet been considered for advertisement to this neighbor.

Task ID

Task ID	Operations
bgp	read

Examples

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

```
RP/0/0/CPU0:router#show bgp summary

BGP router identifier 10.0.0.0, local AS number 2
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 1
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process         RecvTblVer      bRIB/RIB   LabelVer   ImportVer   SendTblVer
Speaker                  1             0          1           1            0

Neighbor        Spk     AS MsgRcvd MsgSent    TblVer   InQ OutQ Up/Down   St/PfxRcd
10.0.101.0        0      2       0       0         0     0    0 00:00:00  Idle
10.0.101.1        0      2       0       0         0     0    0 00:00:00  Idle
```

This table describes the significant fields shown in the display.

Table 36: show bgp summary Field Descriptions

Field	Description
BGP router identifier	IP address of the router.
local AS number	Autonomous system number set by the router bgp, page 231 command. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.

Field	Description
	• 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.
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.

Field	Description
Spr	Speaker process that is responsible for the neighbor. Always 0.
AS	Autonomous system.
MsgRcvd	Number of BGP messages received from a neighbor.
MsgSent	Number of BGP messages sent to a neighbor.
TblVer	Last version of the BGP database that was sent to a neighbor.
InQ	Number of messages from a neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to a neighbor.
Up/Down	Length of time in (hh:mm:ss) that the BGP session has been in Established state, or the time since the session left Established state, if it is not established.
St/PfxRcd	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.
	If the number of prefixes received exceeds the maximum allowed (as set by the **maximum-prefix** command), "(PfxRcd)" appears.
	If the connection has been shut down using the **shutdown** command, "(Admin)" appears.
	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.
	If the connection has been shut down due to out of memory (OOM), "(OOM)" appears.

Related Commands

Command	Description
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.

Command	Description
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

show bgp summary nsr

To display the summary of Border Gateway Protocol (BGP) neighbor state and nonstop routing (NSR) state information, use the **show bgp summary nsr** command in EXEC mode.

show bgp summary [**ipv4** {**unicast**| **multicast**| **labeled-unicast**| **all**| **tunnel**| **mdt**}| **ipv6** {**unicast**| **multicast**| **all**| **labeled-unicast**}| **all** {**unicast**| **multicast**| **all**| **labeled-unicast**| **mdt**| **tunnel**}| **vpnv4 unicast**| **vrf** {*vrf-name*| **all**} [**ipv4** {**unicast**| **labeled-unicast**}| **ipv6 unicast**]| **vpvn6 unicast**] **nsr** [**standby**]

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
ipv6	(Optional) Specifies IP Version 6 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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
standby	Displays information about the standby card.

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

Task ID

Task ID	Operations
bgp	read

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

```
RP/0/0/CPU0:router# show bgp summary nsr

BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 11034 (Reached)
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


node0_1_CPU0        Speaker

Entered mode  Standby Ready               : Feb  3 14:22:00
Entered mode  TCP NSR Setup               : Feb  3 14:22:00
Entered mode  TCP NSR Setup Done          : Feb  3 14:22:01
Entered mode  TCP Initial Sync            : Feb  3 14:22:01
Entered mode  TCP Initial Sync Done       : Feb  3 14:22:44
Entered mode  FPBSN processing done       : Feb  3 14:22:44
Entered mode  Update processing done      : Feb  3 14:22:44
Entered mode  BGP Initial Sync            : Feb  3 14:22:44
Entered mode  BGP Initial Sync done       : Feb  3 14:22:49
Entered mode  NSR Ready                   : Feb  3 14:22:49

Current BGP NSR state - NSR Ready achieved at: Feb  3 14:22:49
NSR State READY notified to Redcon at: Feb  4 07:44:43

Process      RcvTblVer    bRIB/RIB    LabelVer   ImportVer  SendTblVer  StandbyVer
Speaker         13037       13037       13037       13037       13037       13037

Neighbor      Spk    AS    TblVer  SyncVer    AckVer NBRState    NSRState
2.2.2.2        0    302    13037    13037    13037 Established  NSR Ready
```

```
10.0.101.5        0    100    13037    13037    13037 Established  NSR Ready
```

The following example shows sample output from the **show bgp summary nsr** command with the **standby** keyword:

```
RP/0/0/CPU0:router# show bgp summary nsr standby

BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 0 (Not Reached)
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


node0_0_CPU0        Speaker

Entered mode  Standby Ready              : Feb  3 14:22:03
Entered mode  TCP Replication            : Feb  3 14:22:03
Entered mode  TCP Init Sync Done         : Feb  3 14:22:47
Entered mode  NSR Ready                  : Feb  3 14:22:52

Process       RcvTblVer   bRIB/RIB   LabelVer   ImportVer   SendTblVer   StandbyVer
Speaker          13037          0          0       13037            0            0

Neighbor      Spk    AS   TblVer  SyncVer   AckVer NBRState      NSRState
2.2.2.2         0   302    13037        0        1 Established   NSR Ready
10.0.101.5      0   100    13037        0        1 Established   NSR Ready
```

This table describes the significant fields shown in the display.

Table 37: show bgp summary nsr Field Descriptions

Field	Description
BGP router identifier	IP address of the router.
BGP generic scan interval	Interval (in seconds) between scans of the BGP table by a generic scanner.
Non-stop routing	State of the Nonstop routing.
BGP table state	State of the BGP database.
Table ID	BGP database identifier.
BGP main routing table version	Last version of the BGP database that was injected into the main routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
BGP is operating in	Specifies BGP is operating in standalone mode.
Entered mode	The successive transition of various states of TCP and BGP, leading to the NSR ready state.

Field	Description
	Note This is used for monitoring and debugging purposes.
SyncVer	The version which has synced to standby for this neighbor.
AckVer	The version which the neighbor has acknowledge.
NBRState	State of the BGP neighbor.
NSRState	Neighbor NSR state.

Related Commands

Command	Description
nsr (BGP), page 197	Activates Border Gateway Protocol (BGP) nonstop routing (NSR)
show bgp nsr, page 341	Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information.

show bgp table

To display neighbors in the global table, use the **show bgp table** command in EXEC mode.

show bgp table [**ipv4** {**unicast**| **multicast**| **tunnel**| **mdt**}| **ipv6** {**unicast**| **multicast**}| **vpnv4 unicast**| **vpnv6 unicast**]

Syntax Description

ipv4 unicast	(Optional) Specifies IP Version 4 (IPv4) unicast neighbors.
ipv4 multicast	(Optional) Specifies IPv4 multicast neighbors.
ipv4 tunnel	(Optional) Specifies IPv4 tunnel neighbors.
ipv4 mdt	(Optional) Specifies IPv4 multicast distribution tree (MDT) neighbors.
ipv6 unicast	(Optional) Specifies IP Version 6 (IPv6) unicast neighbors.
ipv6 multicast	(Optional) Specifies IPv6 multicast neighbors.
vpnv4 unicast	(Optional) Specifies VPN Version 4 (VPNv4) unicast neighbors.
vpnv6 unicast	(Optional) Specifies VPN Version 6 (VPNv6) unicast neighbors.

Command Default No default behavior or value

Command Modes EXEC

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

The **show bgp table** command provides information about the BGP neighbors based on the global address family and independent of VRFs. Use the **show bgp table** command to get information about all BGP neighbors in different address families.

Task ID

Task ID	Operations
bgp	read

Examples   The following is sample output from the **bgp table vpnv4 unicast** command in EXEC mode:

```
RP/0/0/CPU0:router# show bgp table vpnv4 unicast
Thu Jan 15 17:43:31.215 UTC
Neighbor        VRF              Spk  AS  TblVer  InQ OutQ  St/PfxRcd
10.0.101.1      default           0   1     951   0   0          11
10.0.101.2      default           0   1     951   0   0           5
10.0.101.3      default           0   1     951   0   0           0
10.0.101.4      default           0   1     951   0   0           0
10.0.101.5      default           0   1     951   0   0           0
10.0.101.6      default           0   1     951   0   0           0
10.0.101.7      default           0   1     951   0   0           0
10.0.101.8      default           0   1     951   0   0           0
10.0.101.9      default           0   1     951   0   0           0
90.0.0.2        900               0   2     951   0   0           1
91.0.0.2        901               0   2     951   0   0           1
92.0.0.2        902               0   2     951   0   0           1
93.0.0.2        903               0   2     951   0   0           3
94.0.0.2        904               0   2     951   0   0           3
95.0.0.2        905               0   2     951   0   0           3
96.0.0.2        906               0   2     951   0   0           3
97.0.0.2        907               0   2     951   0   0           3
98.0.0.2        908               0   2     951   0   0           3
99.0.0.2        909               0   2       0   0   0 Idle
12.13.14.16     red               0   2       0   0   0 Idle
20.0.101.1      red               0   2       0   0   0 Active
1.2.3.4         this-is-a-long-vrf-name
                                  0   5       0   0   0 Idle
```

This table describes the significant fields shown in the display.

Table 38: show bgp table Field Descriptions

Field	Description
Neighbor	IP address of a neighbor.
VRF	The VRF which each neighbor belongs to; either the default VRF or a specified VRF.
Spk	Speaker process that is responsible for the neighbor. Always 0.
AS	Autonomous system.
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.
St/PfxRcd	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.

Field	Description
	If the number of prefixes received exceeds the maximum allowed (as set by the maximum-prefix command), "(PfxRcd)" appears.
	If the connection has been shut down using the shutdown command, "(Admin)" appears.
	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.
	If the connection has been shut down due to out of memory (OOM), "(OOM)" appears.

Related Commands

Command	Description
show bgp neighbor-group, page 306	Displays information about the Border Gateway Protocol (BGP) configuration for neighbor groups.
show bgp neighbors, page 310	Displays information about Border Gateway Protocol (BGP) connections to neighbors.
show bgp summary, page 396	Displays the status of all Border Gateway Protocol (BGP) connections.

show bgp truncated-communities

To display routes in the Border Gateway Protocol (BGP) routing table for which inbound policy or aggregation has exceeded the maximum number of communities that may be attached, use the **show bgp truncated-communities** command in EXEC mode.

show bgptruncated-communities standby

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.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
multicast	(Optional) Specifies multicast address prefixes.
ipv6	(Optional) Specifies IP Version 6 address prefixes.
all	(Optional) For address family, specifies prefixes for all address families.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpvn6 unicast	(Optional) Specifies VPNv6 unicast address families.
standby	(Optional) Displays information about the standby card.

Command Default   If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes   EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
	The **count-only** keyword was added.
Release 3.3.0	The following keywords and arguments were added:
	• **vrf** { *vrf-name* | **all** }
	• [**ipv4** { **unicast** | **labeled-unicast** }]
	• [**vpnv4** **unicast**] [**rd** *rd-address*]
Release 3.4.0	The **labeled-unicast** keyword was supported.
	The **count-only** keyword was removed.
Release 3.5.0	The **vpnv6 unicast** [**rd** *rd-address*] keywords and argument were added.
	The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families.
	The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families.
	The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined.

If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

Use the **show bgp truncated-communities** command to display those routes in the specified BGP routing table in which the buffers used to store communities or extended communities have overflowed. An overflow occurs if an attempt is made to associate more communities or extended communities with the route than fits in a BGP update message. This can happen due to modification of communities or extended communities during aggregation or when inbound policy is applied.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp truncated-communities** command:

```
RP/0/0/CPU0:router# show bgp truncated-communities

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 3042
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*  10.13.0.0/16     192.168.40.24                         0 1878 704 701 200 ?
*> 10.16.0.0/16     192.168.40.24                         0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 39: show bgp truncated-communities Field Descriptions

Field	Description
BGP router identifier	BGP Identifier for the local system.
local AS number	Autonomous system number for the local system.
BGP main routing table version	Last version of the BGP database that was installed into the main routing table.
Dampening enabled	Displayed if dampening is enabled for the routes in this BGP routing table.
BGP scan interval	Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family.
Status codes	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):

Field	Description
	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.
	The second character may be (in order of precedence):
	>—Path is the best path to use for that network.
	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.
	The third character may be:
	i—Path was learned by an internal BGP (iBGP) session.
Origin codes	Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:
	i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** 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.
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.

Field	Description
Weight	Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor.
Path	Autonomous system path to the destination network. At the end of the path is the origin code for the path.

Related Commands

Command	Description
aggregate-address, page 24	Creates an aggregate entry in a BGP routing table.
network (BGP), page 184	Specifies a local network that the BGP routing process should originate and advertise to its neighbors.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.
show bgp, page 244	Displays entries in the BGP routing table.

show bgp update-group

To display Border Gateway Protocol (BGP) information for update groups, use the **show bgp update-group** command in EXEC mode.

show bgp [ipv4 {unicast| multicast| labeled-unicast| all| tunnel| mdt}| ipv6 {unicast| multicast| all| labeled-unicast}| all {unicast| multicast| all| labeled-unicast| mdt| tunnel}| vpnv4 unicast| vrf {vrf-name| all} [ipv4 {unicast| labeled-unicast}| ipv6 unicast]| vpnv6 unicast] update-group [neighbor ip-address| process-id.index [summary| performance-statistics]] [standby]

Syntax Description

ipv4	(Optional) Specifies IP Version 4 update groups.
unicast	(Optional) Specifies unicast update groups.
multicast	(Optional) Specifies multicast update groups.
labeled-unicast	(Optional) Specifies labeled unicast address prefixes.
all	(Optional) Displays both unicast and multicast update groups.
tunnel	(Optional) Specifies tunnel address prefixes.
mdt	(Optional) Specifies multicast distribution tree (MDT) address prefixes.
ipv6	(Optional) Specifies IP Version 6 update groups.
all	(Optional) Displays both IP Version 4 and IP Version 6 update groups.
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.
vrf-name	(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.
ipv6 unicast	(Optional) For VRF, specifies IPv6 unicast address families.
vpnv6 unicast	(Optional) Specifies VPNv6 unicast address families.
neighbor *ip-address*	(Optional) Specifies information on an update group for a specific neighbor.
process-id.index	(Optional) Update group index. Process ID range is 0 to 254. Index range is 0 to 4294967295.

The *process id.index* argument is specified as follows: process ID (dot) index. In standalone mode, the process ID is always 0.

summary	(Optional) Specifies summary of update group members.
performance-statistics	(Optional) Specifies performance information about the updates generated for the update group.
standby	(Optional) Displays information about the standby card.

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	The following keywords and argument were added: • **vrf** { *vrf-name* | **all** } • [**ipv4** { **unicast** | **labeled-unicast** }] • [**vpnv4 unicast**]
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **vpnv6 unicast** keywords were added. The **tunnel** and **mdt** keywords were supported under the **ipv4** and **all** address families. The **labeled-unicast** keyword was supported under the **ipv6** and **all** address families. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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. See the *Cisco IOS XR System Management Command Reference for the Cisco XR 12000 Series Router* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Every BGP neighbor is automatically assigned to an update group for each address family that is enabled on the neighbor. Neighbors that have similar outbound policy, such that they are sent the same updates, are placed in the same update group.

Use the **show bgp update-group** command to display the update groups and a list of the neighbors that belong to the update group.

Use the **show bgp update-group neighbor** command to display details about the update group to which a neighbor belongs for the specified address family.

Use the **summary** keyword to display a summary of the neighbors belonging to the specified update group. The display format is the same as for the command.

Use the **performance-statistics** keyword to display information about the number of prefixes processed and the time taken to generate updates for the specified update group.

> **Note** Update group indexes are not necessarily persistent over a process restart. If a BGP process restarts, the index of the update group to which a particular neighbor is assigned may be different, though the set of neighbors belonging to the update group is the same.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp update-group** command:

```
RP/0/0/CPU0:router# show bgp update-group

Update group for IPv4 Unicast, index 0.1:
  Attributes:
    Internal
    Common admin
    Send communities
    Send extended communities
    Minimum advertisement interval: 300
  Update group desynchronized: 0
  Sub-groups merged: 0
  Messages formatted: 0, replicated: 0
  Neighbors not in any sub-group:
    10.0.101.1
```

This table describes the significant fields shown in the display.

Table 40: show bgp update-group Field Descriptions

Field	Description
Update group for	Address family to which updates in this update group apply.
index	Update group index.
Attributes	Attributes common to all members of the update group.
Unsuppress map	Unsuppress route map used to selectively unsuppress more specific routes of locally generated aggregates for members of this update group.
Outbound policy	Route policy applied to outbound updates generated for members of this update group.
Internal	Members of the update group are internal peers.
ORF Receive enabled	Members of this update group are capable of receiving an outbound route filter.
Route Reflector Client	Local system is acting as a route reflector for members of this update group.
Remove private AS numbers	Members of this update group have private AS numbers stripped from outbound updates.
Next-hop-self enabled	Next- Next hop for members of the update group is set to the local router.
Directly connected IPv6 EBGP	Members of this update group are directly connected external BGP IPv6-based peers.
Configured Local AS	Local autonomous system (AS) used for members of this update group.
Common admin	Peers in this update group are under common administration (internal or confederation peers).
Send communities	Communities are sent to neighbors in this update group.
Send extended communities	Extended communities is sent to neighbors in this update group.
Minimum advertisement interval	Minimum advertisement interval for members of this update group.

Field	Description
replicated	Number of update messages replicated for this update group.
Messages formatted	Number of update messages generated for this update group.
Neighbors in this update group	List of neighbors that use this update group for the given address family.
Update group desynchronized	Number of times an update group has been split to accommodate the slower peer. This option is disabled.
Sub-groups merged	Number of times an update group has been split and merged.
Neighbors not in any sub-group	BGP neighbor that does not belong to any subgroup.

The following is sample output from the **show bgp update-group** command with the **ipv4** , **unicast** , and **summary** keywords and the *process id.index* argument:

```
RP/0/0/CPU0:router# show bgp ipv4 unicast update-group 0.1 summary

BGP router identifier 10.140.140.1, local AS number 1.1
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        Spr    AS MsgRcvd MsgSent    TblVer  InQ OutQ Up/Down  St/PfxRcd

172.25.11.8      0     1      0       0         0    0    0   00:00:00 Idle
```

This table describes the significant fields shown in the display.

Table 41: show bgp ipv4 unicast update-group Field Descriptions

Field	Description
BGP router identifier	IP address of the router.
local AS number	Autonomous system number set by the router bgp, page 231command. • 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.

Field	Description
	• Range for 4-byte Autonomous system numbers (ASNs) is 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	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 that exist 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 or advertised to those neighbors.
Neighbor	IP address of a neighbor.
Spr	Speaker process that is responsible for the neighbor. Always 0.

Field	Description
AS	Autonomous system.
MsgRcvd	Number of BGP messages received from a neighbor.
MsgSent	Number of BGP messages sent to a neighbor.
TblVer	Last version of the BGP database that was sent to a neighbor.
InQ	Number of messages from a neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to a neighbor.
Up/Down	Length of time (in hh:mm:s) 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	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. If the number of prefixes received exceeds the maximum allowed (as set by the **maximum-prefix** command), "(PfxRcd)" appears. If the connection has been shut down using the **shutdown** command, "(Admin)" appears. 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.

Related Commands

Command	Description
maximum-prefix (BGP), page 174	Limits the number of prefixes that can be received from a neighbor.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.
set default-afi	Sets the default Address Family Identifier (AFI) for the current session.
set default-safi	Sets the default Subaddress Family Identifier (SAFI) for the current session.

Command	Description
show bgp summary, page 396	Displays the status of all BGP connections.
shutdown (BGP), page 431	Disables a neighbor without removing its configuration.

show bgp vrf imported-routes

To display Border Gateway Protocol (BGP) information for routes imported into specified VPN routing and forwarding (VRF) instances, use the **show bgp vrf imported-routes** command in EXEC mode.

show bgp vrf{**vrf**| *vrf-name*} **imported-routes**

Syntax Description

vrf-name	Displays imported routes for a specific VRF.
all	Displays imported routes for all VRFs.
ipv4 { **unicast** | **labeled-unicast** }	(Optional) Specifies IP Version 4 unicast or labeled-unicast imported routes.
ipv6 unicast	(Optional) Specifies IP Version 6 unicast imported routes.
vrf *source-vrf-name*	(Optional) Displays routes imported from the specified source VRF.
neighbor *neighbor-address*	(Optional) Displays preview advertisements for a specified neighbor.
standby	(Optional) Displays information about the standby card.

Command Default   No default behavior or values

Command Modes   EXEC

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.4.0	The **labeled-unicast** keyword was supported.
Release 3.5.0	The **ipv6 unicast** keywords were added. The **standby** keyword was removed.
Release 3.8.0	The **standby** keyword was added.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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 vrf imported-routes** command to display all paths imported into a specified VRF from the default VRF. Use the **neighbor** *neighbor-address* keyword and argument to display all imported paths and which paths were learned from the specified neighbor. Use the **vrf** *source-vrf-name* keyword and argument to display all imported routes that belong to the specified source VRF. The **neighbor** *neighbor-address* and **vrf** *source-vrf-name* cannot coexist.

Task ID

Task ID	Operations
bgp	read

Examples

The following is sample output from the **show bgp vrf imported-routes** command:

```
RP/0/0/CPU0:router# show bgp vrf vrf-1 ipv6 unicast imported-routes

BGP VRF one, state: Active BGP
BGP Route Distinguisher: 100:222
VRF ID: 0x60000001
BGP router identifier 10.2.0.1, local AS number 100
BGP table state: Active
Table ID: 0xe0800001
BGP main routing table version 41534

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Neighbor       Route Distinguisher    Source VRF
*>i1234:1052::/32    10.1.0.1        100:111                default
*>i2008:1:1:1::/112  10.1.0.1        100:111                default
*>i2008:111:1:1::1/128
                     10.1.0.1        100:111                default

Processed 3 prefixes, 3 paths
```

This table describes the significant fields shown in the display.

Table 42: show bgp vrf imported-routes Field Descriptions

Field	Description
BGP VRF	VRF name.
state	State of the VRF.
BGP Route Distinguisher:	Unique identifier for the BGP routing instance.
VRF Id	VRF identifier.
BGP router identifier	IP address of the router.

Field	Description
local AS number	Autonomous system number set by the router bgp, page 231 command. • 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.
BGP table state	State of the BGP database.
Table ID	Table identifier.
BGP main routing table version	Last version of the BGP database that was injected into the main routing table.
Network	Network address.
Neighbor	IP address of a neighbor.
Route Distinguisher	Unique identifier for the routing instance.
Source VRF	Source VRF for the imported route.

show protocols (BGP)

To display information about the Border Gateway Protocol (BGP) instances running on the router, use the **show protocols** command in EXEC mode and specify either the **bgp** or **all** keyword.

show protocols [**ipv4**| **ipv6**| **afi-all**] [**all**| *protocol*]

Syntax Description

ipv4	(Optional) Specifies the IP Version 4 address family.
ipv6	(Optional) Specifies the IP Version 6 address family.
afi-all	(Optional) Specifies all address families.
all	(Optional) Specifies all protocols for a given address family.
protocol	(Optional) Specifies a routing protocol. For the IPv4 address family, the options are **bgp** , **isis** , **rip** , **eigrp** , and **ospf** . For the IPv6 address family, the options are **bgp** , **eigrp** , **isis** , and **ospfv3** .

Command Default

Default is IPv4.

Command Modes

EXEC

Command History

Release	Modification
Release 3.2	This command was introduced..
Release 3.3.0	The **afi-all** keyword was added.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations.

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 protocols** command to get information about the protocols running on the router and to quickly determine which protocols are active. The command is designed to summarize the important characteristics of the running protocol, and command output varies depending on the specific protocol selected. For BGP, the command output lists the protocol ID, peers with elapsed time since last reset, and miscellaneous information, such as external and internal local distances and sourced routes.

Task ID

Task ID	Operations
bgp	read
rib	read

Examples

The following example shows the display for the **show protocols** command using the **bgp** keyword:

```
RP/0/0/CPU0:router# show protocols bgp

Routing Protocol "BGP 40"

Address Family IPv4 Unicast:
  Distance: external 20 internal 200 local 200
  Sourced Networks:
    10.100.0.0/16 backdoor
    10.100.1.0/24
    10.100.2.0/24
  Routing Information Sources:
    Neighbor        State/Last update received
    10.5.0.2        Idle
    10.9.0.3        Idle
```

This table describes the significant fields shown in the display.

Table 43: show protocols (BGP) Field Descriptions

Field	Description
Routing Protocol:	Identifies BGP as the running protocol and displays the BGP AS number. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535.
Address Family	Specifies the address family. This can be IPv4 Unicast, IPv4 Multicast, or IPv6 Unicast.
Distance: external	Specifies the distance BGP sets when installing eBGP routes into the RIB. eBGP routes are routes received from eBGP peers. The RIB uses the distance as a tiebreaker when several protocols install a route for the same prefix.
Distance: internal	Specifies the distance BGP sets for routes received from iBGP peers.

Field	Description
Distance: local	Specifies the distance BGP sets for locally generated aggregates and backdoor routes.
Sourced Networks	List of locally sourced networks. These are networks sourced using the **network** command.
Routing information Sources	List of configured BGP neighbors.
Neighbor	Address of a BGP neighbor.
State/Last update received	State of each neighbor and the time since the last update was received from the neighbor if it is established.

show svd role

To display Selective VRF Download (SVD) role information, use the **show svd role** command in EXEC mode.

show svd role

Syntax Description   This command has no arguments or keywords.

Command Default   None.

Command Modes   EXEC

Command History

Release	Modification
Release 4.1.0	This command was introduced.

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

The **show svd role** command output displays name of the line card and role for each AFI in a table.

Task ID

Task ID	Operation
ip-services	read

Examples   This example displays the different nodes in a line card and corresponding IPv4, and IPv6 SVD role information:

```
RP/0/0/CPU0:router#show svd role
Thu Mar 10 10:45:17.886 PST
Node Name      IPv4 Role           IPv6 Role
-------------------------------------------------
0/1/CPU0       Core Facing         Not Interested
0/2/CPU0       Core Facing         Core Facing
0/4/CPU0       Standard            Standard
0/5/CPU0       Standard            Standard
```

Related Commands

Command	Description
selective-vrf-download disable, page 233	Disables selective VRF Download (SVD).

Command	Description
show svd state, page 430	Displays Selective VRF Download (SVD) state information.

show svd state

To display Selective VRF Download (SVD) state information, use the show svd state command in EXEC mode.

show svd state

Syntax Description   This command has no arguments or keywords.

Command Default   None.

Command Modes   EXEC

Command History

Release	Modification
Release 4.1.0	This command was introduced.

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

Task ID

Task ID	Operation
ip-services	read

Examples   This example shows the SVD configuration state and the SVD operational state in a line card:

```
RP/0/0/CPU0:router#show svd state
Thu Mar 10 10:45:32.184 PST
Selective VRF Download (SVD) Feature State:
SVD Configuration State      Enabled
SVD Operational State        Enabled
```

Related Commands

Command	Description
selective-vrf-download disable, page 233	Disables selective VRF Download (SVD).
show svd role, page 428	Displays Selective VRF Download (SVD) role information.

shutdown (BGP)

To disable a neighbor without removing its configuration, use the **shutdown** command in an appropriate configuration mode. To re-enable the neighbor and reestablish a Border Gateway Protocol (BGP) session, use the **no** form of this command.

shutdown [inheritance-disable]

no shutdown [inheritance-disable]

Syntax Description		
inheritance-disable	(Optional) Overrides the value of a **shutdown** command inherited from a neighbor group or session group.	

Command Default   Neighbors are not shutdown.

Command Modes   Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the VRF neighbor configuration mode.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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 **shutdown** command to terminate any active session for the specified neighbor and remove all associated routing information. Use of the **shutdown** command with a neighbor group or session group may suddenly terminate a large number of BGP neighbor sessions because all neighbors using the neighbor group or session group may be affected.

Use the **show bgp summary** command to display a summary of BGP neighbors. Neighbors that are idle due to the **shutdown** command are displayed with the "Idle (Admin)" state.

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 that any active session for neighbor 192.168.40.24 is disabled:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# shutdown
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

In the following example, the session remains active for neighbor 192.168.40.24 because the inherited **shutdown** command has been overridden:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# session-group group1
RP/0/0/CPU0:router(config-bgp-sngrp)# shutdown
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/0/CPU0:router(config-bgp-nbr)# shutdown inheritance-disable
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.
show bgp summary, page 396	Displays the status of all BGP connections.

site-of-origin (BGP)

To attach a site-of-origin extended community attribute to each route received from the specified peer, use the **site-of-origin** command in VRF neighbor address family configuration mode. To restore the system to its default condition, use the **no** form of this command.

site-of-origin [*as-number:nn*| *ip-address:nn*]

Syntax Description

as-number:nn	• *as-number—* Autonomous system (AS) number.
	
	  ◦ Range for 2-byte Autonomous system number is 1 to 65535.
	
	  ◦ Range for 4-byte Autonomous system number in asplain format is 1 to 4294967295.
	
	  ◦ Range for 4-byte Autonomous system number is asdot format is 1.0 to 65535.6553.
	
	 • *nn* —32-bit number

ip-address:nn	IP address. • *ip-address* —32-bit IP address • *nn* —16-bit number

Command Default

No default behavior or values

Command Modes

VRF neighbor address family configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.9.0	Asplain format for 4-byte Autonomous system numbers notation was supported.

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.

When routes are advertised to the peer, routes whose extended communities list contain the site of origin (SoO) are filtered out and not advertised to the peer. Site-of-origin uniquely identifies the site from which the

provide edge (PE) router learned routes, thus filtering based on the extended community helps prevent transient routing loops from occurring in complex and mixed network topologies.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure SoO filtering:

```
RP/0/0/CPU0:router(config)# router bgp 6
RP/0/0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# site-of-origin 10.0.01:20
```

socket receive-buffer-size

To set the size of the receive buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket receive-buffer-size** command in an appropriate configuration mode. To set the size of the receive buffers to the default size, use the **no** form of this command.

socket receive-buffer-size *socket-size* [*bgp-size*]

no socket receive-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size (in bytes) of the receive-side socket buffers. Range is 512 to 131072.
bgp-size	(Optional) Size (in bytes) of the receive buffers in BGP. Range is 512 to 131072.

Command Default

socket-size : 32,768 bytes

bgp-size : 4,032 bytes

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the VRF configuration mode.

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 **socket receive-buffer-size** command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.

Note   Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **receive-buffer-size** command on individual neighbors to change the values set by the **socket receive-buffer-size** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the receive buffer sizes for all neighbors to 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# socket receive-buffer-size 65536 8192
```

Related Commands

Command	Description
receive-buffer-size, page 209	Sets the size of the receive buffers for a BGP neighbor.
socket send-buffer-size, page 437	Sets the size of the send buffers for all BGP neighbors.

socket send-buffer-size

To set the size of the send buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default size, use the **no** form of this command.

socket send-buffer-size *socket-size* [*bgp-size*]

no socket send-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

socket-size	Size (in bytes) of the send-side socket buffers. Range is 4096 to 131072.
bgp-size	(Optional) Size (in bytes) of the send buffers in BGP. Range is 4096 to 131072.

Command Default

socket-size : 10240 bytes

bgp-size : 4096 bytes

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the VRF configuration mode.

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 **socket send-buffer-size** command to increase the buffer size when sending updates to neighbors. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.

Note Increasing the socket buffer size uses more memory only when more messages are waiting to be sent by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **send-buffer-size** command on individual neighbors to change the values set by the **socket send-buffer-size** command.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to set the send buffer sizes for all neighbors to 8192 bytes for the socket buffer and the BGP buffer:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# socket send-buffer-size 8192 8192
```

Related Commands

Command	Description
send-buffer-size, page 234	Sets the size of the send buffers for a BGP neighbor.
socket receive-buffer-size, page 435	Sets the size of the receive buffers for all BGP neighbors.

soft-reconfiguration inbound

To configure the software to store updates received from a neighbor, use the **soft-reconfiguration inbound** command in an appropriate configuration mode. To disable storing received updates, use the **no** form of this command.

soft-reconfiguration inbound [**always**| **inheritance-disable**]

no soft-reconfiguration inbound [**always**| **inheritance-disable**]

Syntax Description

always	(Optional) Always performs a soft inbound clear using stored updates, even if the neighbor supports the route refresh capability.
inheritance-disable	(Optional) Overrides configuration for this command that may be inherited from a neighbor group or address family group.

Command Default   Soft reconfiguration is not enabled.

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

VRF IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the following configuration modes: • VPNv4 address family group • VPNv4 neighbor address family

Release	Modification
	• VRF IPv4 neighbor address family
	• VPNv4 neighbor group address family
Release 3.5.0	This command was supported in the following configuration modes:
	• VPNv6 address family group
	• VPNv6 neighbor address family
	• VRF IPv6 neighbor address family
	• VPNv6 neighbor group address family
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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.

To filter or modify some of the updates received from a neighbor, you configure an inbound policy using the **route-policy (BGP)** command. Configuring soft reconfiguration inbound causes the software to store the original unmodified route beside a route that is modified or filtered. This allows a "soft clear" to be performed after the inbound policy is changed. To perform a soft clear, use the **clear bgp soft** command with the **in** keyword specified. The unmodified routes are then passed through the new policy and installed in the BGP table.

Note If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors.

Note The bgp auto-policy-soft-reset is enabled by default. A soft clear is done automatically when the inbound policy configured with the **route-policy** (BGP) command is changed. This behavior can be changed by disabling the auto-policy-soft-reset using the **bgp auto-policy-soft-reset disable** command.

If the neighbor supports the route refresh capability, then the original routes are not stored because they can be retrieved from the neighbor through a route refresh request. However, if the **always** keyword is specified, the original routes are stored even when the neighbor supports the route refresh capability.

If the **soft-reconfiguration inbound** command is not configured and the neighbor does not support the route refresh capability, then an inbound soft clear is not possible. In that case, the only way to rerun the inbound policy is to use the **clear bgp** *ip-address* command to reset the neighbor BGP session.

Note If there is an existing BGP session with a neighbor that does not support the route refresh capability, the session is terminated and a new one is initiated.

✎

Note   The extra routes stored as a result of configuring this command use more memory on the router.

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

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows inbound soft reconfiguration enabled for IP Version 4 (IPv4) unicast routes received from neighbor 10.108.1.1. The software stores all routes received in their unmodified form so that when an inbound soft clear is performed later, the stored information can then be used to generate a new set of modified routes.

```
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# neighbor 10.108.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound
RP/0/0/CPU0:router(config-bgp-nbr-af)# exit
```

The following example shows inbound soft reconfiguration disabled for neighbor 10.108.1.1, preventing this feature from being automatically inherited by address family group group1:

```
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# soft-reconfiguration inbound
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.108.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound inheritance-disable
RP/0/0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
bgp auto-policy-soft-reset disable, page 42	Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified.
clear bgp, page 104	Resets a BGP connection using a soft or hard reset.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.

Command	Description
rd, page 207	Applies a prefix list to filter updates received from a neighbor.
route-policy (BGP), page 225	Applies a routing policy to updates advertised to or received from a BGP neighbor.

speaker-id

To allocate a speaker process to a neighbor, use the **speaker-id** command in the appropriate configuration mode. To remove the speaker process from a neighbor, use the **no** form of this command.

speaker-id *id*

no speaker-id [*id*]

Syntax Description

id	ID of the speaker process. Range is 1 to 15.

Command Default

Default is 0.

Command Modes

Neighbor configuration

Session group configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.6.0	The command was supported in session group configuration mode.

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.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to allocate speaker process 3 to neighbor 192.168.40.24:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# speaker-id 3
```

Related Commands

Command	Description
distributed speaker, page 141	Starts a specified speaker process.

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

policy-name	Name of the routing policy to apply.

Command Default

No policy is applied when routes are installed into the routing table.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VRF IPv4 address family configuration

VRF IPv6 address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the VRF IPv4 address family configuration mode.
Release 3.5.0	This command was supported in the VRF IPv6 address family configuration mode.

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 a routing 'black hole' where BGP advertises routes to neighbors that BGP does not install in its global routing table and forwarding table.

Use the **table-policy** command to modify route attributes as the routes are installed into the routing table by Border Gateway Protocol (BGP). Commonly, it is used to set the traffic index attribute.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to apply the set-traffic-index policy to IPv4 unicast routes being installed into the routing table:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-af)# table-policy set-traffic-index
```

Related Commands

Command	Description
route-policy (RPL)	Defines a route policy and enters route policy configuration mode.

timers (BGP)

To set the timers for a specific Border Gateway Protocol (BGP) neighbor, use the **timers** command in an appropriate configuration mode. To set the timers to the default values, use the **no** form of this command.

timers *keepalive hold-time*

no timers [*keepalive hold-time*]

Syntax Description

keepalive	Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 4294967295.
hold-time	Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 4294967295.

Command Default

keepalive : 60 seconds

hold-time : 180 seconds

Use the **timers bgp** command to override the default values.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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 timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

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

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbours that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

Bear in mind that BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to change the keepalive timer to 70 seconds and the hold-time timer to 210 seconds for the BGP peer 192.168.40.24:

```
RP/0/0/CPU0:router(config)# router bgp 109
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# timers 70 210
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.
timers bgp, page 448	Adjusts BGP network timers for all BGP neighbors.

timers bgp

To change the default timer values for Border Gateway Protocol (BGP) neighbors, use the **timers bgp** command in an appropriate configuration mode. To set the default timers to the default values, use the **no** form of this command.

timers bgp *keepalive hold-time*

no timers bgp [*keepalive hold-time*]

Syntax Description

keepalive	Frequency (in seconds) with which the software sends keepalive messages to the neighbor. Range is 0 to 65535.
hold-time	Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 65535.

Command Default

keepalive : 60 seconds

hold-time : 180 seconds

Command Modes

Router configuration

VRF configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF configuration mode.

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 **timers bgp** command to adjust the default timer times used by all BGP neighbors. The values can be overridden on particular neighbors using the **timers** command in the neighbor configuration mode.

The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbours that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

Bear in mind that BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to configure a default keepalive time of 30 seconds and a default hold time of 90 seconds:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# timers bgp 30 90
```

Related Commands

Command	Description
timers (BGP), page 446	Adjusts BGP network timers for a BGP neighbor.

ttl-security

To configure a router to check the time-to-live (TTL) field in incoming IP packets for the specified external Border Gateway Protocol (eBGP) peer, use the **ttl-security** command in an appropriate configuration mode. To disable TTL verification, use the **no** form of this command.

ttl-security [inheritance-disable]

no ttl-security [inheritance-disable]

Syntax Description

inheritance-disable	(Optional) Prevents the **ttl-security** command from being inherited from a session group or neighbor group.

Command Default

TTL verification is not enabled for eBGP peers.

Command Modes

Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.
Release 3.9.0	The **disable** keyword was replaced with the **inheritance-disable** keyword.

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 **ttl-security** command to enable a lightweight security mechanism to protect eBGP peering sessions from CPU utilization-based and other resource exhaustion-based attacks. These types of attacks are typically brute-force Denial of Service (DoS) attacks that attempt to disable the network by flooding devices in the network with IP packets that contain forged source and destination IP addresses in the packet headers.

This command leverages existing behavior in IP packets. For a given IP packet, the TTL count of the packet always is equal to or less than the TTL count when the packet originated, a behavior that is considered impossible to circumvent. Therefore, a packet received with a TTL count equal to the maximum TTL value of 255 can be sent only by a directly adjacent peer. When the **ttl-security** command is configured for an eBGP neighbor that is directly adjacent, the router accepts only IP packets with a TTL count that is equal to the maximum TTL value.

The **ttl-security** command secures the eBGP session in the incoming direction only. In the outbound direction, it causes packets to be sent only with the maximum TTL value so that the BGP neighbor can also verify the TTL value of incoming packets. When this command is enabled, BGP establishes or maintains a session only if the TTL value in the IP packet header is equal to the maximum TTL value. If the value is less than the maximum TTL value, the packet is discarded and an Internet Control Message Protocol (ICMP) message is not generated. This behavior is designed because a response to a forged packet is not necessary.

Note   The **ttl-security** command must be configured on each participating router. Failure to configure this command on both ends of the BGP session results in the session progressing as far as the OpenSent or OpenConfirm state, remaining there until the hold time expires.

The following restrictions apply to the configuration of this command:

- The **ttl-security** command should not be configured for a peer that is already configured with the **neighbor ebgp-multihop** command. The simultaneous configuration of these commands is permitted; however, the **ttl-security** command overrides the **ebgp-multihop** command.

- This command is not supported for internal BGP (iBGP) peers.

- This command is not effective against attacks from a directly adjacent peer that has been compromised.

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

Note   If the **ttl-security** command is configured on a neighbor to which the router has an established connection or the router is in the process of establishing a connection, the session must be cleared using the **clear bgp** command.

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to enable TTL security for eBGP neighbor 192.168.223.7:

```
RP/0/0/CPU0:router(config)# router bgp 65534
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.223.7
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 65507
RP/0/0/CPU0:router(config-bgp-nbr)# ttl-security
```

The following example shows how to enable TTL security for multiple eBGP neighbors using a session group:

```
RP/0/0/CPU0:router(config)# router bgp 65534
RP/0/0/CPU0:router(config-bgp)# session-group ebgp-nbrs
RP/0/0/CPU0:router(config-bgp-sngrp)# ttl-security
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.223.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 65501
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/0/CPU0:router(config-bgp-nbr)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.223.2
```

```
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 65502
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/0/CPU0:router(config-bgp-nbr)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.223.3
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 65503
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
ebgp-multihop, page 146	Accepts and attempts BGP connections to external peers residing on networks that are not directly connected.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.
show lpts flows	Displays information about locally terminated packet flows, including the minimum TTL value expected.

update-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

type	Interface type. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface. **Note** Use the **show interfaces** command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (**?**) online help function.

Command Default   Best local address

Command Modes   Neighbor configuration

VRF neighbor configuration

Neighbor group configuration

Session group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.
Release 3.5.0	No modification

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/0/0/CPU0:router(config)# router bgp 110
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 110
RP/0/0/CPU0:router(config-bgp-nbr)# update-source Loopback0
```

Related Commands

Command	Description
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.

use

To inherit configuration from a neighbor group, session group, or address family group, use the **use** command in an appropriate configuration mode. To discontinue inheritance from a group, use the **no** form of this command.

use {**af-group** *group-name*| **neighbor-group** *group-name*| **session-group** *group-nam* **e**}

no use {**af-group** [*group-name*]| **neighbor-group** [*group-name*]| **session-group** [*group-name*]}

Syntax Description

af-group	Specifies an address family group.
group-name	Name of the neighbor group, session group, or address family group from which you want to inherit configuration.
neighbor-group	Specifies a neighbor group.
session-group	Specifies a session group.

Command Default   Inheritance of group characteristics does not occur.

Command Modes   For **use af-group** version:

Address family group configuration

Neighbor address family configuration

Neighbor group address family configuration

For **use neighbor-group** version:

Neighbor group configuration

Neighbor configuration

VRF neighbor configuration

For **use session-group** version:

Neighbor group configuration

Neighbor configuration

VRF neighbor configuration

Session-group configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in VRF neighbor configuration mode.

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 **use** command configures inheritance of configuration from an address family group, neighbor group, or session group, which means that any configuration for the group also takes effect for the user of the group.

The configuration inherited depends on the type of group that is specified. The group types are described in the following sections:

Address Family Group

An address family group can specify a configuration for only a single address family. The address family specified when the address family group was defined (through the **af-group** command) must match the address family from which the group is used.

Neighbor Group

A neighbor group (like a neighbor) can have address family-independent configuration and address family-specific configuration. All of these configurations could be inherited.

Session Group

A session group can have only address family-independent configuration and thus only address family-independent configuration is inherited from it.

The following rules govern inheritance to resolve possible conflicting configuration:

1   If a command is configured directly on the neighbor that is using group configuration, the command overrides the value that would be normally inherited from the group.
2   If the neighbor is configured to use a session group (for address family-independent configuration) or an address family group (for address family-specific configuration) and the command is configured for the session group or address family group, that configuration is used.
3   The neighbor group configuration is used:

   • If the command is not configured directly on the neighbor and the neighbor is not using a session group (for address family-independent configuration) or an af-group (for address family-specific configuration).

   • The neighbor is using a neighbor group and the command is configured on the neighbor group.

Typically, all configuration for a neighbor group is inherited, but some characteristics may be masked by a session group or address family group. For an example of this configuration, see the "Examples" section.

If the neighbor is using both a session group and a neighbor group and a specific command is configured for the neighbor group but not for the session group, then the configuration for the neighbor group does not take effect. The session group "hides" all address family-independent configuration on the neighbor group and prevents it from being inherited. Similarly, the use of an address family group hides any address family-specific configuration that may otherwise be inherited from a neighbor group for that address family.

In addition to neighbors using groups, it is possible to build a hierarchy by having groups use other groups. The following hierarchical groups are permitted:

   • Session groups may use other session groups.

   • Address family groups may use other address family groups.

   • Neighbor groups may use other neighbor groups.

> • Neighbor groups may use session groups and address family groups.

Note　Within the Cisco IOS XR system configuration architecture, do not combine the **remote-as** command and the **no use neighbor-group** command in the same commit, or the **remote-as** command and the **no use session-group** command in the same commit.

Task ID

Task ID	Operations
bgp	read, write

Examples　The following example shows how to define a session group session1 and configure neighbor 172.168.40.24 to use session1. As a result, the session1 configuration takes effect on the neighbor also.

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# session-group session1
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

The following example is similar to the previous example, but in this case the **timers** command on the session group does not take effect on the neighbor because it is overridden by a **timers** command directly configured for the neighbor.

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# session-group session1
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/0/CPU0:router(config-bgp-nbr)# timers 60 180
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

The following example shows an address family group, family1, for IPv4 multicast and a neighbor group, neighbor1, that have IPv4 unicast and IPv4 multicast enabled. In this case, the neighbor inherits IPv4 unicast (and address family-independent) configuration from the neighbor group, but inherits IPv4 multicast configuration from the address family group. In this example, the neighbor group also has a remote autonomous system configured, so there is no need to configure a remote autonomous system for the neighbor because it inherits the remote autonomous system from the neighbor group:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# af-group family1 address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-afgrp)# route-policy mcast-in in
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor-group neighbor1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# remote-as 2
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 in
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 out
```

```
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 in
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 out
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group neighbor1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# use af-group family1
RP/0/0/CPU0:router(config-bgp-nbr-af)# exit
```

In the previous example, the neighbor uses the policy1 route policy for inbound and outbound IPv4 unicast routes, but uses the mcast-in route policy for inbound IPv4 multicast routes and no policy for outbound IPv4 multicast routes.

The following example shows a neighbor inheriting configuration from a session group that likewise inherits configuration from another session group. The configuration from both session groups take effect on the neighbor:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# session-group session1
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# session-group session2
RP/0/0/CPU0:router(config-bgp-sngrp)# use session-group session1
RP/0/0/CPU0:router(config-bgp-sngrp)# update-source Loopback0
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group session2
RP/0/0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
remote-as (BGP), page 215	Creates a BGP neighbor and begins the exchange of routing information.
show bgp af-group, page 264	Displays information about BGP configuration for address family groups.
show bgp neighbor-group, page 306	Displays information about the BGP configuration for neighbor groups.
show bgp neighbors, page 310	Displays information about BGP neighbors.
show bgp session-group, page 390	Displays information about the BGP configuration for session groups.

vrf (BGP)

To configure a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove the VRF instance from the configuration file and restore the system to its default condition, use the **no** form of this command.

vrf *vrf-name*

no vrf *vrf-name*

Syntax Description

vrf-name	Name of the VRF instance. The following names cannot be used: all, default, and global.

Command Default　　No default behavior or values

Command Modes　　Router configuration

Command History

Release	Modification
Release 3.3.0	This command was introduced.
Release 3.5.0	The following restriction was removed:
	If you remove a VRF configuration using the **no vrf** *vrf-name* command and want to reconfigure the VRF configuration using the **vrf** *vrf-name* command, you must wait at least three minutes.

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 **vrf** command to configure a VRF instance. A VRF instance is a collection of VPN routing and forwarding tables maintained at the provider edge (PE) router.

Task ID

Task ID	Operations
bgp	read, write

Examples　　The following example shows how to configure a VRF instance and enter VRF configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/0/CPU0:router(config-bgp-vrf)#
```

weight

To assign a weight to routes received from a neighbor, use the **weight** command in an appropriate configuration mode. To remove the **weight** command from the configuration file and restore the system to its default condition in which the software assigns the default weight to routes, use the **no** form of this command.

weight *weight-value*

no weight [*weight-value*]

Syntax Description

weight-value	Weight to assign. Range is 0 to 65535.

Command Default

Routes learned through another Border Gateway Protocol (BGP) peer have a default weight of 0 and routes sourced by the local router have a default weight of 32768.

Command Modes

IPv4 address family group configuration

IPv6 address family group configuration

VPNv4 address family group configuration

IPv4 neighbor address family configuration

VPNv4 neighbor address family configuration

VRF IPv4 neighbor address family configuration

IPv4 neighbor group address family configuration

IPv6 neighbor group address family configuration

VPNv4 neighbor group address family configuration

VPNv6 address family group configuration

VPNv6 neighbor address family configuration

VRF IPv6 neighbor address family configuration

VPNv6 neighbor group address family configuration

Command History

Release	Modification
Release 3.2	This command was introduced.
Release 3.3.0	This command was supported in the following configuration modes: • VPNv4 address family group • VPNv4 neighbor address family • VRF IPv4 neighbor address family • VPNv4 neighbor group address family

Release	Modification
Release 3.5.0	This command was supported in the following configuration modes:
	• VPNv6 address family group configuration
	• VPNv6 neighbor address family configuration
	• VRF IPv6 neighbor address family configuration
	• VPNv6 neighbor group address family configuration

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 weight of a route is a Cisco-specific attribute. It is used in the best-path selection process (as the strongest tie-breaker). See the *Implementing BGP on Cisco IOS XR Software* module of the *Cisco IOS XR Routing Configuration Guide for the Cisco XR 12000 Series Router* for information on best path. If there are two BGP routes with the same network layer reachability information (NLRI), the route with the higher weight is always chosen no matter what the value of other BGP attributes. Weight only has significance on the local router. Weight is assigned locally to the router, is a value that only makes sense to the specific router, is not propagated or carried through any route updates, and never is sent between BGP peers (even within the same AS).

Note If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors.

The weight assigned to individual routes can be further manipulated in the inbound route policy of a neighbor using the **set weight** command. The **set weight** command sets the weight directly. If you have particular neighbors that you want to prefer for most of your outbound traffic, you can assign a higher weight to all routes learned from that neighbor.

The weight assigned to individual routes may be modified by using an inbound routing policy.

Note For weight changes to take effect, you may need to use the clear bgp soft, page 128 command.

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

Task ID

Task ID	Operations
bgp	read, write

Examples

The following example shows how to assign a weight of 50 to all IP Version 4 (IPv4) unicast routes learned through 172.20.16.6:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# neighbor 172.20.16.6
```

```
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbr-af)# weight 50
RP/0/0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

Command	Description
af-group, page 22	Creates an address family group for BGP neighbors and enters address family group configuration mode.
clear bgp, page 104	Resets a group of BGP neighbors.
neighbor-group, page 182	Creates a neighbor group and enters neighbor group configuration mode.
session-group, page 240	Creates a session group and enters session group configuration mode.
set weight	Sets the weight for BGP routes.

weight reset-on-import

To reset weight of paths on import, use the **weight reset-on-import** command in an appropriate configuration mode. To remove the **weight reset-on-import** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

weight reset-on-import

no weight reset-on-import

Syntax Description   This command has no arguments or keywords.

Command Default   Reset weight on import is disabed.

Command Modes   VRF IPv4 address family configuration

VRF IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to reset weight of paths on import under VRF IPv4 address family configuration:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# vrf vrf1
RP/0/0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-vrf-af)# weight reset-on-import
```

The following example shows how to reset weight of paths on import under VPNv6 address family configuration:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# address-family vpnv6 unicast
RP/0/0/CPU0:router(config-bgp-af)# weight reset-on-import
```

Related Commands

Command	Description
weight reset-on-import disable, page 465	Disables resetting weight of paths on import, if it is enabled globally.

weight reset-on-import disable

To disable resetting weight of paths on import, if it is enabled globally, use the **weight reset-on-import-disable** in appropriate configuration mode. To cancel the disable option and retain the weight reset-on-import option globally, use the **no** form of this command.

weight reset-on-import disable

no weight reset-on-import disable

Syntax Description   This command has no arguments or keywords.

Command Default   Reset weight of paths on import option is enabled globally.

Command Modes   VRF IPv4 address family configuration

VRF IPv6 address family configuration

VPNv4 address family configuration

VPNv6 address family configuration

Command History

Release	Modification
Release 3.8.0	This command was introduced.

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

Task ID

Task ID	Operations
bgp	read, write

Examples   The following example shows how to disable reset weight of paths on import option under VPNv4 address family configuration:

```
RP/0/0/CPU0:router(config)# router bgp 1
RP/0/0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-vrf-af)# weight reset-on-import disable
```

Related Commands

Command	Description
weight reset-on-import, page 463	Reset weight of paths on import.