



# Implementing BGP on Cisco IOS XR Software

Border Gateway Protocol (BGP) is an Exterior Gateway Protocol (EGP) that allows you to create loop-free interdomain routing between autonomous systems. An *autonomous system* is a set of routers under a single technical administration. Routers in an autonomous system can use multiple Interior Gateway Protocols (IGPs) to exchange routing information inside the autonomous system and an EGP to route packets outside the autonomous system.

This module provides the conceptual and configuration information for BGP on Cisco IOS XR software.



## Note

For more information about BGP on the Cisco IOS XR software and complete descriptions of the BGP commands listed in this module, see [Related Documents](#), [page 158](#) section of this module. To locate documentation for other commands that might appear while performing a configuration task, search online in the Cisco IOS XR software master command index.

## Feature History for Implementing BGP

Release	Modification
Release 3.2	This feature was introduced.
Release 3.3.0	VPN routing and forwarding (VRF) support was added, including information on VRF command modes and command syntax. BGP cost community information was added.
Release 3.4.0	The following features were supported: <ul style="list-style-type: none"><li>• Four-byte autonomous system (AS) number</li><li>• Carrier supporting carrier (CSC) for BGP was added. See <i>Cisco IOS XR Multiprotocol Label Switching Protocol Configuration Guide</i> for information</li><li>• Key chains</li></ul>
Release 3.5.0	The following features were supported:

Release	Modification
	<ul style="list-style-type: none"> <li>• IPv6 Provider Edge and IPv6 VPN Provider Edge over Multiprotocol Label Switching</li> <li>• Neighbor-specific VRF IPv6 address family configurations</li> <li>• Address family group-specific VPNv6 configurations</li> <li>• VPN4/VPNv6 over IP core using L2TPv3 tunnels</li> <li>• Multicast Distribution Tree (MDT) Subaddress Family Identifier Information (SAFI) support for multicast VPN (MVPN)</li> </ul>
Release 3.6.0	No modification.
Release 3.7.0	<p>The following features were supported:</p> <ul style="list-style-type: none"> <li>• Advertisement of VRF routes for multicast VPNs (MVPN) for both IPv4 and IPv6 address families from PE to PE</li> <li>• Edits were made to existing MVPN procedures based on new support for IPv6 multicast VPNs</li> <li>• Procedure <a href="#">Configuring an MDT Address Family Session in BGP, page 131</a> was updated to reflect MVPN configuration of MDT SAFI from PE to PE</li> </ul>
Release 3.8.0	<p>The following features were supported:</p> <ul style="list-style-type: none"> <li>• Border Gateway Protocol (BGP) nonstop routing (NSR) with stateful switchover (SSO)</li> <li>• Next hop as the IPv6 address of peering interface</li> <li>• Reset weight on import of VPN routes</li> <li>• New commands <b>enforce-first-as</b> and <b>enforce-first-as-disable</b> were introduced to provide enable and disable configuration options for enforce-first-as feature in Neighbor, Neighbor group, and Session group configuration modes.</li> </ul>
Release 3.9.0	<p>The following features were supported:</p> <ul style="list-style-type: none"> <li>• BGP Best-External Path</li> <li>• BGP Prefix Independent Convergence Unipath Primary Backup</li> <li>• BGP Local Label Retention</li> <li>• Asplain notation for 4-byte Autonomous System Number</li> </ul>

Release	Modification
	<ul style="list-style-type: none"> <li>• Command Line Interface (CLI) consistency for BGP commands</li> <li>• L2VPN Address Family Configuration Mode</li> </ul>

- [Prerequisites for Implementing BGP, page 3](#)
- [Information About Implementing BGP, page 3](#)
- [How to Implement BGP on Cisco IOS XR Software, page 49](#)
- [Configuration Examples for Implementing BGP, page 152](#)
- [Where to Go Next, page 157](#)
- [Additional References, page 158](#)

## Prerequisites for Implementing BGP

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

## Information About Implementing BGP

To implement BGP, you need to understand the following concepts:

### BGP Functional Overview

BGP uses TCP as its transport protocol. Two BGP routers form a TCP connection between one another (peer routers) and exchange messages to open and confirm the connection parameters.

BGP routers exchange network reachability information. This information is mainly an indication of the full paths (BGP autonomous system numbers) that a route should take to reach the destination network. This information helps construct a graph that shows which autonomous systems are loop free and where routing policies can be applied to enforce restrictions on routing behavior.

Any two routers forming a TCP connection to exchange BGP routing information are called peers or neighbors. BGP peers initially exchange their full BGP routing tables. After this exchange, incremental updates are sent as the routing table changes. BGP keeps a version number of the BGP table, which is the same for all of its BGP peers. The version number changes whenever BGP updates the table due to routing information changes. Keepalive packets are sent to ensure that the connection is alive between the BGP peers and notification packets are sent in response to error or special conditions.

**Note**

---

For information on configuring BGP to distribute Multiprotocol Label Switching (MPLS) Layer 3 virtual private network (VPN) information, see the Cisco IOS XR Multiprotocol Label Switching Configuration Guide for the Cisco XR 12000 Series Router

For information on BGP support for Bidirectional Forwarding Detection (BFD), see the *Cisco IOS XR Interface and Hardware Configuration Guide for the Cisco XR 12000 Series Router* and the *Cisco IOS XR Interface and Hardware Command Reference for the Cisco XR 12000 Series Router*.

---

## BGP Router Identifier

For BGP sessions between neighbors to be established, BGP must be assigned a router ID. The router ID is sent to BGP peers in the OPEN message when a BGP session is established.

BGP attempts to obtain a router ID in the following ways (in order of preference):

- By means of the address configured using the **bgp router-id** command in router configuration mode.
- By using the highest IPv4 address on a loopback interface in the system if the router is booted with saved loopback address configuration.
- By using the primary IPv4 address of the first loopback address that gets configured if there are not any in the saved configuration.

If none of these methods for obtaining a router ID succeeds, BGP does not have a router ID and cannot establish any peering sessions with BGP neighbors. In such an instance, an error message is entered in the system log, and the **show bgp summary** command displays a router ID of 0.0.0.0.

After BGP has obtained a router ID, it continues to use it even if a better router ID becomes available. This usage avoids unnecessary flapping for all BGP sessions. However, if the router ID currently in use becomes invalid (because the interface goes down or its configuration is changed), BGP selects a new router ID (using the rules described) and all established peering sessions are reset.

**Note**

---

We strongly recommend that the **bgp router-id** command is configured to prevent unnecessary changes to the router ID (and consequent flapping of BGP sessions).

---

## BGP Default Limits

Cisco IOS XR BGP imposes maximum limits on the number of neighbors that can be configured on the router and on the maximum number of prefixes that are accepted from a peer for a given address family. This limitation safeguards the router from resource depletion caused by misconfiguration, either locally or on the remote neighbor. The following limits apply to BGP configurations:

- The default maximum number of peers that can be configured is 4000. The default can be changed using the **bgp maximum neighbor** command. The *limit* range is 1 to 15000. Any attempt to configure additional peers beyond the maximum limit or set the maximum limit to a number that is less than the number of peers currently configured will fail.
- To prevent a peer from flooding BGP with advertisements, a limit is placed on the number of prefixes that are accepted from a peer for each supported address family. The default limits can be overridden

through configuration of the maximum-prefix *limit* command for the peer for the appropriate address family. The following default limits are used if the user does not configure the maximum number of prefixes for the address family:

- 512K (524,288) prefixes for IPv4 unicast.
- 128K (131,072) prefixes for IPv4 multicast.
- 128K (131,072) prefixes for IPv6 unicast.
- 128K (131,072) prefixes for IPv6 multicast
- 512K (524,288) prefixes for VPNv4 unicast
- 512K (524,288) prefixes for VPNv6 unicast

A cease notification message is sent to the neighbor and the peering with the neighbor is terminated when the number of prefixes received from the peer for a given address family exceeds the maximum limit (either set by default or configured by the user) for that address family.

It is possible that the maximum number of prefixes for a neighbor for a given address family has been configured after the peering with the neighbor has been established and a certain number of prefixes have already been received from the neighbor for that address family. A cease notification message is sent to the neighbor and peering with the neighbor is terminated immediately after the configuration if the configured maximum number of prefixes is fewer than the number of prefixes that have already been received from the neighbor for the address family.

## BGP Next Hop Tracking

BGP receives notifications from the Routing Information Base (RIB) when next-hop information changes (event-driven notifications). BGP obtains next-hop information from the RIB to:

- Determine whether a next hop is reachable.
- Find the fully recursed IGP metric to the next hop (used in the best-path calculation).
- Validate the received next hops.
- Calculate the outgoing next hops.
- Verify the reachability and connectedness of neighbors.

BGP is notified when any of the following events occurs:

- Next hop becomes unreachable
- Next hop becomes reachable
- Fully recursed IGP metric to the next hop changes
- First hop IP address or first hop interface change
- Next hop becomes connected
- Next hop becomes unconnected
- Next hop becomes a local address
- Next hop becomes a nonlocal address

**Note**

Reachability and recurred metric events trigger a best-path recalculation.

Event notifications from the RIB are classified as critical and noncritical. Notifications for critical and noncritical events are sent in separate batches. However, a noncritical event is sent along with the critical events if the noncritical event is pending and there is a request to read the critical events.

- Critical events are related to the reachability (reachable and unreachable), connectivity (connected and unconnected), and locality (local and nonlocal) of the next hops. Notifications for these events are not delayed.
- Noncritical events include only the IGP metric changes. These events are sent at an interval of 3 seconds. A metric change event is batched and sent 3 seconds after the last one was sent.

The next-hop trigger delay for critical and noncritical events can be configured to specify a minimum batching interval for critical and noncritical events using the **nexthop trigger-delay** command. The trigger delay is address family dependent.

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 route that resolved a next hop as well as the mask length associated with the route. The **nexthop route-policy** command is used to specify the route-policy.

For information on route policy filtering for next hops using the next-hop attach point, see the *Implementing Routing Policy Language on Cisco IOS XR Software* module of *Cisco IOS XR Routing Configuration Guide* (this publication).

## Next Hop as the IPv6 Address of Peering Interface

BGP can carry IPv6 prefixes over an IPv4 session. The next hop for the IPv6 prefixes can be set through a nexthop policy. In the event that the policy is not configured, the nexthops are set as the IPv6 address of the peering interface (IPv6 neighbor interface or IPv6 update source interface, if any one of the interfaces is configured).

If the nexthop policy is not configured and neither the IPv6 neighbor interface nor the IPv6 update source interface is configured, the next hop is the IPv4 mapped IPv6 address.

## Scoped IPv4/VPNv4 Table Walk

To determine which address family to process, a next-hop notification is received by first dereferencing the gateway context associated with the next hop, then looking into the gateway context to determine which address families are using the gateway context. The IPv4 unicast and VPNv4 unicast address families share the same gateway context, because they are registered with the IPv4 unicast table in the RIB. As a result, both the global IPv4 unicast table and the VPNv4 table are processed when an IPv4 unicast next-hop notification is received from the RIB. A mask is maintained in the next hop, indicating whether the next hop belongs to IPv4 unicast or VPNv4 unicast, or both. This scoped table walk localizes the processing in the appropriate address family table.

## Reordered Address Family Processing

The Cisco IOS XR software walks address family tables based on the numeric value of the address family. When a next-hop notification batch is received, the order of address family processing is reordered to the following order:

- IPv4 tunnel
- VPNv4 unicast
- VPNv6 unicast
- IPv4 labeled unicast
- IPv4 unicast
- IPv4 MDT
- IPv4 multicast
- IPv6 unicast
- IPv6 multicast
- IPv6 labeled unicast

## New Thread for Next-Hop Processing

The critical-event thread in the spkr process handles only next-hop, Bidirectional Forwarding Detection (BFD), and fast-external-failover (FEF) notifications. This critical-event thread ensures that BGP convergence is not adversely impacted by other events that may take a significant amount of time.

## show, clear, and debug Commands

The **show bgp nexthops** command provides statistical information about next-hop notifications, the amount of time spent in processing those notifications, and details about each next hop registered with the RIB. The **clear bgp nexthop performance-statistics** command ensures that the cumulative statistics associated with the processing part of the next-hop **show** command can be cleared to help in monitoring. The **clear bgp nexthop registration** command performs an asynchronous registration of the next hop with the RIB. See the *BGP Commands on Cisco IOS XR Software* module of *Cisco IOS XR Routing Command Reference for the Cisco XR 12000 Series Router* for information on the next-hop **show** and **clear** commands.

The **debug bgp nexthop** command displays information on next-hop processing. The **out** keyword provides debug information only about BGP registration of next hops with RIB. The **in** keyword displays debug information about next-hop notifications received from RIB. The **out** keyword displays debug information about next-hop notifications sent to the RIB. See the *BGP Debug Commands on Cisco IOS XR Software* module of *Cisco IOS XR Routing Debug Command Reference for the Cisco XR 12000 Series Router* .

## Autonomous System Number Formats in BGP

Autonomous system numbers (ASNs) are globally unique identifiers used to identify autonomous systems (ASs) and enable ASs to exchange exterior routing information between neighboring ASs. A unique ASN is allocated to each AS for use in BGP routing. ASNs are encoded as 2-byte numbers and 4-byte numbers in BGP.

## 2-byte Autonomous System Number Format

The 2-byte ASNs are represented in asplain notation. The 2-byte range is 1 to 65535.

## 4-byte Autonomous System Number Format

To prepare for the eventual exhaustion of 2-byte Autonomous System Numbers (ASNs), BGP has the capability to support 4-byte ASNs. The 4-byte ASNs are represented both in asplain and asdot notations.

The byte range for 4-byte ASNs in asplain notation is 1-4294967295. The AS is represented as a 4-byte decimal number. The 4-byte ASN asplain representation is defined in draft-ietf-idr-as-representation-01.txt.

For 4-byte ASNs in asdot format, the 4-byte range is 1.0 to 65535.65535 and the format is:

*high-order-16-bit-value-in-decimal . low-order-16-bit-value-in-decimal*

The BGP 4-byte ASN capability is used to propagate 4-byte-based AS path information across BGP speakers that do not support 4-byte AS numbers. See draft-ietf-idr-as4bytes-12.txt for information on increasing the size of an ASN from 2 bytes to 4 bytes. AS is represented as a 4-byte decimal number

## as-format Command

The **as-format** command configures the ASN notation to asdot. The default value, if the **as-format** command is not configured, is asplain.

## BGP Configuration

BGP in Cisco IOS XR software follows a neighbor-based configuration model that requires that all configurations for a particular neighbor be grouped in one place under the neighbor configuration. Peer groups are not supported for either sharing configuration between neighbors or for sharing update messages. The concept of peer group has been replaced by a set of configuration groups to be used as templates in BGP configuration and automatically generated update groups to share update messages between neighbors.

## Configuration Modes

BGP configurations are grouped into modes. The following sections show how to enter some of the BGP configuration modes. From a mode, you can enter the **?** command to display the commands available in that mode.

### Router Configuration Mode

The following example shows how to enter router configuration mode:

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

### Router Address Family Configuration Mode

The following example shows how to enter router address family configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 112
```



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

## Neighbor Configuration Mode

The following example shows how to enter neighbor configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/0/CPU0:router(config-bgp-nbr)#
```

## Neighbor Address Family Configuration Mode

The following example shows how to enter neighbor address family configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 112
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbr-af)#
```

## VRF Configuration Mode

The following example shows how to enter VPN routing and forwarding (VRF) configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/0/CPU0:router(config-bgp-vrf)#
```

## VRF Address Family Configuration Mode

The following example shows how to enter VRF address family configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 112
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)#
```

## VRF Neighbor Configuration Mode

The following example shows how to enter VRF neighbor configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/0/CPU0:router(config-bgp-vrf)# neighbor 11.0.1.2
RP/0/0/CPU0:router(config-bgp-vrf-nbr)#
```

## VRF Neighbor Address Family Configuration Mode

The following example shows how to enter VRF neighbor address family configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 112
RP/0/0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/0/CPU0:router(config-bgp-vrf)# neighbor 11.0.1.2
RP/0/0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
```

```
RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)#
```

### VPNv4 Address Family Configuration Mode

The following example shows how to enter VPNv4 address family configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 152
RP/0/0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/0/CPU0:router(config-bgp-af)#
```

### VPNv6 Address Family Configuration Mode

The following example shows how to enter VPNv6 address family configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 150
RP/0/0/CPU0:router(config-bgp)# address-family vpnv6 unicast
RP/0/0/CPU0:router(config-bgp-af)#
```

### L2VPN Address Family Configuration Mode

The following example shows how to enter L2VPN address family configuration mode:

```
RP/0/0/CPU0:router(config)# router bgp 100
RP/0/0/CPU0:router(config-bgp)# address-family l2vpn vpls-vpws
RP/0/0/CPU0:router(config-bgp-af)#
```

## Neighbor Submode

Cisco IOS XR BGP uses a neighbor submode to make it possible to enter configurations without having to prefix every configuration with the **neighbor** keyword and the neighbor address:

- Cisco IOS XR software has a submode available for neighbors in which it is not necessary for every command to have a “neighbor x.x.x.x” prefix:

In Cisco IOS XR software, the configuration is as follows:

```
RP/0/
RP0/CPU0:router(config-bgp)# neighbor 192.23.1.2
RP/0/
RP0/CPU0:router(config-bgp-nbr)# remote-as 2002
RP/0/
RP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
```

- An address family configuration submode inside the neighbor configuration submode is available for entering address family-specific neighbor configurations. In Cisco IOS XR software, the configuration is as follows:

```
RP/0/
RP0/CPU0:router(config-bgp)# neighbor 2002::2
RP/0/
RP0/CPU0:router(config-bgp-nbr)# remote-as 2023
RP/0/
RP0/CPU0:router(config-bgp-nbr)# address-family ipv6 unicast
RP/0/
```

```
RP0/CPU0:router (config-bgp-nbr-af) # next-hop-self
RP/0/
RP0/CPU0:router (config-bgp-nbr-af) # route-policy one in
```

- You must enter neighbor-specific IPv4, IPv6, VPNv4, or VPNv6 commands in neighbor address-family configuration submode. In Cisco IOS XR software, the configuration is as follows:

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

- You must enter neighbor-specific IPv4 and IPv6 commands in VRF neighbor address-family configuration submode. In Cisco IOS XR software, the configuration is as follows:

```
RP/0/
RP0/CPU0:router (config) # router bgp 110
RP/0/
RP0/CPU0:router (config-bgp) # vrf vrf_A
RP/0/
RP0/CPU0:router (config-bgp-vrf) # neighbor 11.0.1.2
RP/0/
RP0/CPU0:router (config-bgp-vrf-nbr) # address-family ipv4 unicast
RP/0/
RP0/CPU0:router (config-bgp-vrf-nbr-af) # route-policy pass all in
```

## Configuration Templates

The **af-group**, **session-group**, and **neighbor-group** configuration commands provide template support for the neighbor configuration in Cisco IOS XR software.

The **af-group** command is used to group address family-specific neighbor commands within an IPv4, IPv6, VPNv4, or VPNv6 address family. Neighbors that have the same address family configuration are able to use the address family group (af-group) name for their address family-specific configuration. A neighbor inherits the configuration from an address family group by way of the **use** command. If a neighbor is configured to use an address family group, the neighbor (by default) inherits the entire configuration from the address family group. However, a neighbor does not inherit all of the configuration from the address family group if items are explicitly configured for the neighbor. The address family group configuration is entered under the BGP router configuration mode. The following example shows how to enter address family group configuration mode :

```
RP/0/0/CPU0:router (config) # router bgp 140
RP/0/0/CPU0:router (config-bgp) # af-group afmcast1 address-family ipv4 multicast
RP/0/0/CPU0:router (config-bgp-afgrp) #
```

The **session-group** command allows you to create a session group from which neighbors can inherit address family-independent configuration. A neighbor inherits the configuration from a session group by way of the **use** command. If a neighbor is configured to use a session group, the neighbor (by default) inherits the entire

configuration of the session group. A neighbor does not inherit all of the configuration from a session group if a configuration is done directly on that neighbor. The following example shows how to enter session group configuration mode:

```
RP/0/0/CPU0:router# router bgp 140
RP/0/0/CPU0:router (config-bgp) # session-group session1
RP/0/0/CPU0:router (config-bgp-sngrp) #
```

The **neighbor-group** command helps you apply the same configuration to one or more neighbors. Neighbor groups can include session groups and address family groups and can comprise the complete configuration for a neighbor. After a neighbor group is configured, a neighbor can inherit the configuration of the group using the **use** command. If a neighbor is configured to use a neighbor group, the neighbor inherits the entire BGP configuration of the neighbor group.

The following example shows how to enter neighbor group configuration mode:

```
RP/0/0/CPU0:router (config) # router bgp 123
RP/0/0/CPU0:router (config-bgp) # neighbor-group nbrgroup1
RP/0/0/CPU0:router (config-bgp-nbrgrp) #
```

The following example shows how to enter neighbor group address family configuration mode:

```
RP/0/0/CPU0:router (config) # router bgp 140
RP/0/0/CPU0:router (config-bgp) # neighbor-group nbrgroup1
RP/0/0/CPU0:router (config-bgp-nbrgrp) # address-family ipv4 unicast
RP/0/0/CPU0:router (config-bgp-nbrgrp-af) #
```

- However, a neighbor does not inherit all of the configuration from the neighbor group if items are explicitly configured for the neighbor. In addition, some part of the configuration of the neighbor group could be hidden if a session group or address family group was also being used.

Configuration grouping has the following effects in Cisco IOS XR software:

- Commands entered at the session group level define address family-independent commands (the same commands as in the neighbor submode).
- Commands entered at the address family group level define address family-dependent commands for a specified address family (the same commands as in the neighbor-address family configuration submode).
- Commands entered at the neighbor group level define address family-independent commands and address family-dependent commands for each address family (the same as all available **neighbor** commands), and define the **use** command for the address family group and session group commands.

## Template Inheritance Rules

In Cisco IOS XR software, BGP neighbors or groups inherit configuration from other configuration groups.

For address family-independent configurations:

- Neighbors can inherit from session groups and neighbor groups.
- Neighbor groups can inherit from session groups and other neighbor groups.
- Session groups can inherit from other session groups.
- If a neighbor uses a session group and a neighbor group, the configurations in the session group are preferred over the global address family configurations in the neighbor group.

For address family-dependent configurations:

- Address family groups can inherit from other address family groups.
- Neighbor groups can inherit from address family groups and other neighbor groups.
- Neighbors can inherit from address family groups and neighbor groups.

Configuration group inheritance rules are numbered in order of precedence as follows:

- 1 If the item is configured directly on the neighbor, that value is used. In the example that follows, the advertisement interval is configured both on the neighbor group and neighbor configuration and the advertisement interval being used is from the neighbor configuration:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 15
RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1
RP/0/0/CPU0:router(config-bgp-nbr)# advertisement-interval 20
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 20 seconds:

```
RP/0/0/CPU0:router# show bgp neighbors 10.1.1.1

BGP neighbor is 10.1.1.1, remote AS 1, local AS 140, 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
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 20 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.1
  eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
  Route refresh request: received 0, sent 0
  0 accepted prefixes
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%

Connections established 0; dropped 0
Last reset 00:00:14, due to BGP neighbor initialized
External BGP neighbor not directly connected.
```

- 2 Otherwise, if an item is configured to be inherited from a session-group or neighbor-group and on the neighbor directly, then the configuration on the neighbor is used. If a neighbor is configured to be inherited from session-group or af-group, but no directly configured value, then the value in the session-group or af-group is used. In the example that follows, the advertisement interval is configured on a neighbor group and a session group and the advertisement interval value being used is from the session group:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# session-group AS_2
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 15
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 20
RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.0.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
```

```
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group AS_2
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 15 seconds:

```
RP/0/0/CPU0:router# show bgp neighbors 192.168.0.1

BGP neighbor is 192.168.0.1, remote AS 1, local AS 140, 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
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 15 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.1
  eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
  Route refresh request: received 0, sent 0
  0 accepted prefixes
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%

Connections established 0; dropped 0
Last reset 00:03:23, due to BGP neighbor initialized
External BGP neighbor not directly connected.
```

- 3 Otherwise, if the neighbor uses a neighbor group and does not use a session group or address family group, the configuration value can be obtained from the neighbor group either directly or through inheritance. In the example that follows, the advertisement interval from the neighbor group is used because it is not configured directly on the neighbor and no session group is used:

```
RP/0/0/CPU0:router(config)# router bgp 150
RP/0/0/CPU0:router(config-bgp)# session-group AS_2
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 20
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 15
RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.1.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 15 seconds:

```
RP/0/0/CPU0:router# show bgp neighbors 192.168.1.1

BGP neighbor is 192.168.2.2, remote AS 1, local AS 140, 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
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 15 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.1
  eBGP neighbor with no outbound policy; defaults to 'drop'
  Route refresh request: received 0, sent 0
  Inbound path policy configured
  Policy for incoming advertisements is POLICY_1
  0 accepted prefixes
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%
```

```

Connections established 0; dropped 0
Last reset 00:01:14, due to BGP neighbor initialized
External BGP neighbor not directly connected.

```

To illustrate the same rule, the following example shows how to set the advertisement interval to 15 (from the session group) and 25 (from the neighbor group). The advertisement interval set in the session group overrides the one set in the neighbor group. The inbound policy is set to POLICY\_1 from the neighbor group.

```

RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# session-group ADV
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 15
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor-group ADV_2
RP/0/0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 25
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# route-policy POLICY_1 in
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)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 192.168.2.2
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbr)# use session-group ADV
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group ADV_2

```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 15 seconds:

```

RP/0/0/CPU0:router# show bgp neighbors 192.168.2.2

BGP neighbor is 192.168.2.2, remote AS 1, local AS 140, 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
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 15 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 0
  Update group: 0.1
  eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
  Route refresh request: received 0, sent 0
  0 accepted prefixes
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%

Connections established 0; dropped 0
Last reset 00:02:03, due to BGP neighbor initialized
External BGP neighbor not directly connected.

```

- 4 Otherwise, the default value is used. In the example that follows, neighbor 10.0.101.5 has the minimum time between advertisement runs set to 30 seconds (default) because the neighbor is not configured to use the neighbor configuration or the neighbor group configuration:

```

RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# neighbor-group AS_1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# remote-as 1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor-group adv_15
RP/0/0/CPU0:router(config-bgp-nbrgrp)# remote-as 10
RP/0/0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 15
RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.101.5
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group AS_1
RP/0/0/CPU0:router(config-bgp-nbr)# exit

```

```
RP/0/0/CPU0:router(config-bgp)# neighbor 10.0.101.10
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group adv_15
```

The following output from the **show bgp neighbors** command shows that the advertisement interval used is 30 seconds:

```
RP/0/0/CPU0:router# show bgp neighbors 10.0.101.5

BGP neighbor is 10.0.101.5, remote AS 1, local AS 140, 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
Received 0 messages, 0 notifications, 0 in queue
Sent 0 messages, 0 notifications, 0 in queue
Minimum time between advertisement runs is 30 seconds

For Address Family: IPv4 Unicast
BGP neighbor version 0
Update group: 0.2
eBGP neighbor with no inbound or outbound policy; defaults to 'drop'
Route refresh request: received 0, sent 0
0 accepted prefixes
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
Connections established 0; dropped 0
Last reset 00:00:25, due to BGP neighbor initialized
External BGP neighbor not directly connected.
```

The inheritance rules used when groups are inheriting configuration from other groups are the same as the rules given for neighbors inheriting from groups.

## Viewing Inherited Configurations

You can use the following **show** commands to view BGP inherited configurations:

### show bgp neighbors

Use the **show bgp neighbors** command to display information about the BGP configuration for neighbors.

- Use the **configuration** keyword to display the effective configuration for the neighbor, including any settings that have been inherited from session groups, neighbor groups, or address family groups used by this neighbor.
- Use the **inheritance** keyword to display the session groups, neighbor groups, and address family groups from which this neighbor is capable of inheriting configuration.

The **show bgp neighbors** command examples that follow are based on this sample configuration:

```
RP/0/0/CPU0:router(config)# router bgp 142
RP/0/0/CPU0:router(config-bgp)# af-group GROUP_3 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# next-hop-self
RP/0/0/CPU0:router(config-bgp-afgrp)# route-policy POLICY_1 in
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# session-group GROUP_2
RP/0/0/CPU0:router(config-bgp-sngrp)# advertisement-interval 15
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor-group GROUP_1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# use session-group GROUP_2
RP/0/0/CPU0:router(config-bgp-nbrgrp)# ebgp-multihop 3
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# weight 100
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)# default-originate
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 192.168.0.1
RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group GROUP_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 GROUP_3
RP/0/0/CPU0:router(config-bgp-nbr-af)# weight 200

```

The following example displays sample output from the **show bgp neighbors** command using the **inheritance** keyword. The example shows that the neighbor inherits session parameters from neighbor group GROUP\_1, which in turn inherits from session group GROUP\_2. The neighbor inherits IPv4 unicast parameters from address family group GROUP\_3 and IPv4 multicast parameters from neighbor group GROUP\_1:

```

RP/0/0/CPU0:router# show bgp neighbors 192.168.0.1 inheritance

Session:          n:GROUP_1 s:GROUP_2
IPv4 Unicast:     a:GROUP_3
IPv4 Multicast:  n:GROUP_1

```

The following example displays sample output from the **show bgp neighbors** command using the **configuration** keyword. The example shows from where each item of configuration was inherited, or if it was configured directly on the neighbor (indicated by [ ]). For example, the **ebgp-multihop 3** command was inherited from neighbor group GROUP\_1 and the **next-hop-self** command was inherited from the address family group GROUP\_3:

```

RP/0/0/CPU0:router# show bgp neighbors 192.168.0.1 configuration

neighbor 192.168.0.1
remote-as 2 [ ]
advertisement-interval 15 [n:GROUP_1 s:GROUP_2]
ebgp-multihop 3 [n:GROUP_1]
address-family ipv4 unicast [ ]
next-hop-self [a:GROUP_3]
route-policy POLICY_1 in [a:GROUP_3]
weight 200 [ ]
address-family ipv4 multicast [n:GROUP_1]
default-originate [n:GROUP_1]

```

## show bgp af-group

Use the **show bgp af-group** command to display address family groups:

- Use the **configuration** keyword to display the effective configuration for the address family group, including any settings that have been inherited from address family groups used by this address family group.
- Use the **inheritance** keyword to display the address family groups from which this address family group is capable of inheriting configuration.
- Use the **users** keyword to display the neighbors, neighbor groups, and address family groups that inherit configuration from this address family group.

The **show bgp af-group** sample commands that follow are based on this sample configuration:

```

RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# af-group GROUP_3 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# remove-private-as
RP/0/0/CPU0:router(config-bgp-afgrp)# route-policy POLICY_1 in

```

```

RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# af-group GROUP_1 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# use af-group GROUP_2
RP/0/0/CPU0:router(config-bgp-afgrp)# maximum-prefix 2500 75 warning-only
RP/0/0/CPU0:router(config-bgp-afgrp)# default-originate
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# af-group GROUP_2 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# use af-group GROUP_3
RP/0/0/CPU0:router(config-bgp-afgrp)# send-community-ebgp
RP/0/0/CPU0:router(config-bgp-afgrp)# send-extended-community-ebgp
RP/0/0/CPU0:router(config-bgp-afgrp)# capability orf prefix both

```

The following example displays sample output from the **show bgp af-group** command using the **configuration** keyword. This example shows from where each configuration item was inherited. The **default-originate** command was configured directly on this address family group (indicated by [ ]). The **remove-private-as** command was inherited from address family group GROUP\_2, which in turn inherited from address family group GROUP\_3:

```

RP/0/0/CPU0:router# show bgp af-group GROUP_1 configuration

af-group GROUP_1 address-family ipv4 unicast
  capability orf prefix-list both           [a:GROUP_2]
  default-originate                         [ ]
  maximum-prefix 2500 75 warning-only      [ ]
  route-policy POLICY_1 in                  [a:GROUP_2 a:GROUP_3]
  remove-private-AS                         [a:GROUP_2 a:GROUP_3]
  send-community-ebgp                       [a:GROUP_2]
  send-extended-community-ebgp             [a:GROUP_2]

```

The following example displays sample output from the **show bgp af-group** command using the **users** keyword:

```

RP/0/0/CPU0:router# show bgp af-group GROUP_2 users

IPv4 Unicast: a:GROUP_1

```

The following example displays sample output from the **show bgp af-group** command using the **inheritance** keyword. This shows that the specified address family group GROUP\_1 directly uses the GROUP\_2 address family group, which in turn uses the GROUP\_3 address family group:

```

RP/0/0/CPU0:router# show bgp af-group GROUP_1 inheritance

IPv4 Unicast: a:GROUP_2 a:GROUP_3

```

## show bgp session-group

Use the **show bgp session-group** command to display session groups:

- Use the **configuration** keyword to display the effective configuration for the session group, including any settings that have been inherited from session groups used by this session group.
- Use the **inheritance** keyword to display the session groups from which this session group is capable of inheriting configuration.
- Use the **users** keyword to display the session groups, neighbor groups, and neighbors that inherit configuration from this session group.

The output from the **show bgp session-group** command is based on the following session group configuration:

```
RP/0/0/CPU0:router(config)# router bgp 113
RP/0/0/CPU0:router(config-bgp)# session-group GROUP_1
RP/0/0/CPU0:router(config-bgp-sngrp)# use session-group GROUP_2
RP/0/0/CPU0:router(config-bgp-sngrp)# update-source Loopback_0
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# session-group GROUP_2
RP/0/0/CPU0:router(config-bgp-sngrp)# use session-group GROUP_3
RP/0/0/CPU0:router(config-bgp-sngrp)# ebgp-multihop 2
RP/0/0/CPU0:router(config-bgp-sngrp)# exit
RP/0/0/CPU0:router(config-bgp)# session-group GROUP_3
RP/0/0/CPU0:router(config-bgp-sngrp)# dmz-link-bandwidth
```

The following is sample output from the **show bgp session-group** command with the **configuration** keyword in EXEC mode:

```
RP/0/0/CPU0:router# show bgp session-group GROUP_1 configuration

session-group GROUP_1
  ebgp-multihop 2          [s:GROUP_2]
  update-source Loopback0 []
  dmz-link-bandwidth      [s:GROUP_2 s:GROUP_3]
```

The following is sample output from the **show bgp session-group** command with the **inheritance** keyword showing that the GROUP\_1 session group inherits session parameters from the GROUP\_3 and GROUP\_2 session groups:

```
RP/0/0/CPU0:router# show bgp session-group GROUP_1 inheritance

Session: s:GROUP_2 s:GROUP_3
```

The following is sample output from the **show bgp session-group** command with the **users** keyword showing that both the GROUP\_1 and GROUP\_2 session groups inherit session parameters from the GROUP\_3 session group:

```
RP/0/0/CPU0:router# show bgp session-group GROUP_3 users

Session: s:GROUP_1 s:GROUP_2
```

## show bgp neighbor-group

Use the **show bgp neighbor-group** command to display neighbor groups:

- Use the **configuration** keyword to display the effective configuration for the neighbor group, including any settings that have been inherited from neighbor groups used by this neighbor group.
- Use the **inheritance** keyword to display the address family groups, session groups, and neighbor groups from which this neighbor group is capable of inheriting configuration.
- Use the **users** keyword to display the neighbors and neighbor groups that inherit configuration from this neighbor group.

The examples are based on the following group configuration:

```
RP/0/0/CPU0:router(config)# router bgp 140
RP/0/0/CPU0:router(config-bgp)# af-group GROUP_3 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# remove-private-as
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)# af-group GROUP_2 address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-afgrp)# use af-group GROUP_3
RP/0/0/CPU0:router(config-bgp-afgrp)# send-community-ebgp
RP/0/0/CPU0:router(config-bgp-afgrp)# send-extended-community-ebgp
RP/0/0/CPU0:router(config-bgp-afgrp)# capability orf prefix both
RP/0/0/CPU0:router(config-bgp-afgrp)# exit
RP/0/0/CPU0:router(config-bgp)# session-group GROUP_3
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-group GROUP_1
RP/0/0/CPU0:router(config-bgp-nbrgrp)# remote-as 1982
RP/0/0/CPU0:router(config-bgp-nbrgrp)# use neighbor-group GROUP_2
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/0/CPU0:router(config-bgp)# neighbor-group GROUP_2
RP/0/0/CPU0:router(config-bgp-nbrgrp)# use session-group GROUP_3
RP/0/0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# use af-group GROUP_2
RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# weight 100

```

The following is sample output from the **show bgp neighbor-group** command with the **configuration** keyword. The configuration setting source is shown to the right of each command. In the output shown previously, the remote autonomous system is configured directly on neighbor group GROUP\_1, and the send community setting is inherited from neighbor group GROUP\_2, which in turn inherits the setting from address family group GROUP\_3:

```
RP/0/0/CPU0:router# show bgp neighbor-group GROUP_1 configuration
```

```

neighbor-group GROUP_1
remote-as 1982                []
timers 30 90                 [n:GROUP_2 s:GROUP_3]
address-family ipv4 unicast  []
capability orf prefix-list both [n:GROUP_2 a:GROUP_2]
remove-private-AS            [n:GROUP_2 a:GROUP_2 a:GROUP_3]
send-community-ebgp          [n:GROUP_2 a:GROUP_2]
send-extended-community-ebgp [n:GROUP_2 a:GROUP_2]
soft-reconfiguration inbound [n:GROUP_2 a:GROUP_2 a:GROUP_3]
weight 100                   [n:GROUP_2]

```

The following is sample output from the **show bgp neighbor-group** command with the **inheritance** keyword. This output shows that the specified neighbor group GROUP\_1 inherits session (address family-independent) configuration parameters from neighbor group GROUP\_2. Neighbor group GROUP\_2 inherits its session parameters from session group GROUP\_3. It also shows that the GROUP\_1 neighbor group inherits IPv4 unicast configuration parameters from the GROUP\_2 neighbor group, which in turn inherits them from the GROUP\_2 address family group, which itself inherits them from the GROUP\_3 address family group:

```
RP/0/0/CPU0:router# show bgp neighbor-group GROUP_1 inheritance
```

```

Session:      n:GROUP-2 s:GROUP_3
IPv4 Unicast: n:GROUP_2 a:GROUP_2 a:GROUP_3

```

The following is sample output from the **show bgp neighbor-group** command with the **users** keyword. This output shows that the GROUP\_1 neighbor group inherits session (address family-independent) configuration parameters from the GROUP\_2 neighbor group. The GROUP\_1 neighbor group also inherits IPv4 unicast configuration parameters from the GROUP\_2 neighbor group:

```
RP/0/0/CPU0:router# show bgp neighbor-group GROUP_2 users
```

```

Session:      n:GROUP_1
IPv4 Unicast: n:GROUP_1

```

## No Default Address Family

BGP does not support the concept of a default address family. An address family must be explicitly configured under the BGP router configuration for the address family to be activated in BGP. Similarly, an address family must be explicitly configured under a neighbor for the BGP session to be activated under that address family. It is not required to have any address family configured under the BGP router configuration level for a neighbor to be configured. However, it is a requirement to have an address family configured at the BGP router configuration level for the address family to be configured under a neighbor.

## Routing Policy Enforcement

External BGP (eBGP) neighbors must have an inbound and outbound policy configured. If no policy is configured, no routes are accepted from the neighbor, nor are any routes advertised to it. This added security measure ensures that routes cannot accidentally be accepted or advertised in the case of a configuration omission error.



### Note

This enforcement affects only eBGP neighbors (neighbors in a different autonomous system than this router). For internal BGP (iBGP) neighbors (neighbors in the same autonomous system), all routes are accepted or advertised if there is no policy.

In the following example, for an eBGP neighbor, if all routes should be accepted and advertised with no modifications, a simple pass-all policy is configured:

```
RP/0/0/CPU0:router(config)# route-policy pass-all
RP/0/0/CPU0:router(config-rpl)# pass
RP/0/0/CPU0:router(config-rpl)# end-policy
RP/0/0/CPU0:router(config)# commit
```

Use the **route-policy (BGP)** command in the neighbor address-family configuration mode to apply the pass-all policy to a neighbor. The following example shows how to allow all IPv4 unicast routes to be received from neighbor 192.168.40.42 and advertise all IPv4 unicast routes back to it:

```
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)# remote-as 21
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 in
RP/0/0/CPU0:router(config-bgp-nbr-af)# route-policy pass-all out
RP/0/0/CPU0:router(config-bgp-nbr-af)# commit
```

Use the **show bgp summary** command to display eBGP neighbors that do not have both an inbound and outbound policy for every active address family. In the following example, such eBGP neighbors are indicated in the output with an exclamation (!) mark:

```
RP/0/0/CPU0:router# show bgp all all summary

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 41
BGP scan interval 60 secs
BGP is operating in STANDALONE mode.
```

```

Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          41           41          41

Neighbor         Spk   AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
10.0.101.1       0     1     919      925      41      0    0  15:15:08  10
10.0.101.2       0     2       0       0       0       0    0  00:00:00  Idle

```

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 1  
BGP scan interval 60 secs  
BGP is operating in STANDALONE mode.

```

Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          1            1            1

```

Some configured eBGP neighbors do not have both inbound and outbound policies configured for IPv4 Multicast address family. These neighbors will default to sending and/or receiving no routes and are marked with '!' in the output below. Use the 'show bgp neighbor <nbr\_address>' command for details.

```

Neighbor         Spk   AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
10.0.101.2       0     2       0       0       0       0    0  00:00:00  Idle!

```

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 2  
BGP scan interval 60 secs  
BGP is operating in STANDALONE mode.

```

Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          2            2            2

```

```

Neighbor         Spk   AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
2222::2          0     2     920      918      2       0    0  15:15:11  1
2222::4          0     3       0       0       0       0    0  00:00:00  Idle

```

Address Family: IPv6 Multicast  
=====

BGP router identifier 10.0.0.1, local AS number 1  
BGP generic scan interval 60 secs  
BGP main routing table version 1  
BGP scan interval 60 secs  
BGP is operating in STANDALONE mode.

```

Process          RecvTblVer    bRIB/RIB    SendTblVer
Speaker          1            1            1

```

Some configured eBGP neighbors do not have both inbound and outbound policies configured for IPv6 Multicast address family. These neighbors will default to sending and/or receiving no routes and are marked with '!' in the output below. Use the 'show bgp neighbor <nbr\_address>' command for details.

```

Neighbor         Spk   AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
2222::2          0     2     920      918      0       0    0  15:15:11  0
2222::4          0     3       0       0       0       0    0  00:00:00  Idle!

```

## Table Policy

The table policy feature in BGP allows you to configure traffic index values on routes as they are installed in the global routing table. This feature is enabled using the **table-policy** command and supports the BGP policy accounting feature.

BGP policy accounting uses traffic indices that are set on BGP routes to track various counters. See the *Implementing Routing Policy on Cisco IOS XR Software* module in the *Cisco IOS XR Routing Configuration Guide for the Cisco XR 12000 Series Router* for details on table policy use. See the *Cisco Express Forwarding Commands on Cisco IOS XR Software* module in the *Cisco IOS XR IP Addresses and Services Command Reference for the Cisco XR 12000 Series Router* for details on BGP policy accounting.

Table policy also provides 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.

## Update Groups

The BGP Update Groups feature contains an algorithm that dynamically calculates and optimizes update groups of neighbors that share outbound policies and can share the update messages. The BGP Update Groups feature separates update group replication from peer group configuration, improving convergence time and flexibility of neighbor configuration.

To use this feature, you must understand the following concepts:

### Related Topics

- [BGP Update Generation and Update Groups](#) , page 23
- [BGP Update Group](#) , page 23

## BGP Update Generation and Update Groups

The BGP Update Groups feature separates BGP update generation from neighbor configuration. The BGP Update Groups feature introduces an algorithm that dynamically calculates BGP update group membership based on outbound routing policies. This feature does not require any configuration by the network operator. Update group-based message generation occurs automatically and independently.

## BGP Update Group

When a change to the configuration occurs, the router automatically recalculates update group memberships and applies the changes.

For the best optimization of BGP update group generation, we recommend that the network operator keeps outbound routing policy the same for neighbors that have similar outbound policies. This feature contains commands for monitoring BGP update groups. For more information about the commands, see [Monitoring BGP Update Groups](#), page 145.

## BGP Cost Community

The BGP cost community is a nontransitive extended community attribute that is passed to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers. The cost community feature allows you to customize the local route preference and influence the best-path selection process by assigning cost values to specific routes. The extended community format defines generic points of insertion (POI) that influence the best-path decision at different points in the best-path algorithm.

The cost community attribute is applied to internal routes by configuring the **set extcommunity cost** command in a route policy. See the *Routing Policy Language Commands on Cisco IOS XR Software* module of *Cisco IOS XR Routing Command Reference* for information on the **set extcommunity cost** command. The cost community set clause is configured with a cost community ID number (0–255) and cost community number (0–4294967295). The cost community number determines the preference for the path. The path with the lowest cost community number is preferred. Paths that are not specifically configured with the cost community number are assigned a default cost community number of 2147483647 (the midpoint between 0 and 4294967295) and evaluated by the best-path selection process accordingly. When two paths have been configured with the same cost community number, the path selection process prefers the path with the lowest cost community ID. The cost-extended community attribute is propagated to iBGP peers when extended community exchange is enabled.

The following commands include the **route-policy** keyword, which you can use to apply a route policy that is configured with the cost community set clause:

- **aggregate-address**
- **redistribute**
- **network**

## How BGP Cost Community Influences the Best Path Selection Process

The cost community attribute influences the BGP best-path selection process at the point of insertion (POI). By default, the POI follows the Interior Gateway Protocol (IGP) metric comparison. When BGP receives multiple paths to the same destination, it uses the best-path selection process to determine which path is the best path. BGP automatically makes the decision and installs the best path in the routing table. The POI allows you to assign a preference to a specific path when multiple equal cost paths are available. If the POI is not valid for local best-path selection, the cost community attribute is silently ignored.

Cost communities are sorted first by POI then by community ID. Multiple paths can be configured with the cost community attribute for the same POI. The path with the lowest cost community ID is considered first. In other words, all cost community paths for a specific POI are considered, starting with the one with the lowest cost community. Paths that do not contain the cost community cost (for the POI and community ID being evaluated) are assigned the default community cost value (2147483647). If the cost community values are equal, then cost community comparison proceeds to the next lowest community ID for this POI.

To select the path with the lower cost community, simultaneously walk through the cost communities of both paths. This is done by maintaining two pointers to the cost community chain, one for each path, and advancing both pointers to the next applicable cost community at each step of the walk for the given POI, in order of community ID, and stop when a best path is chosen or the comparison is a tie. At each step of the walk, the following checks are done:

```
If neither pointer refers to a cost community,  
  Declare a tie;
```



```

Elseif a cost community is found for one path but not for the other,
    Choose the path with cost community as best path;
Elseif the Community ID from one path is less than the other,
    Choose the path with the lesser Community ID as best path;
Elseif the Cost from one path is less than the other,
    Choose the path with the lesser Cost as best path;
Else Continue.

```

**Note**

Paths that are not configured with the cost community attribute are considered by the best-path selection process to have the default cost value (half of the maximum value [4294967295] or 2147483647).

Applying the cost community attribute at the POI allows you to assign a value to a path originated or learned by a peer in any part of the local autonomous system or confederation. The cost community can be used as a “tie breaker” during the best-path selection process. Multiple instances of the cost community can be configured for separate equal cost paths within the same autonomous system or confederation. For example, a lower cost community value can be applied to a specific exit path in a network with multiple equal cost exit points, and the specific exit path is preferred by the BGP best-path selection process. See the scenario described in [Influencing Route Preference in a Multiexit IGP Network](#), page 26.

**Note**

The cost community comparison in BGP is enabled by default. Use the **bgp bestpath cost-community ignore** command to disable the comparison.

See [BGP Best Path Algorithm](#), page 28 for information on the BGP best-path selection process.

## Cost Community Support for Aggregate Routes and Multipaths

The BGP cost community feature supports aggregate routes and multipaths. The cost community attribute can be applied to either type of route. The cost community attribute is passed to the aggregate or multipath route from component routes that carry the cost community attribute. Only unique IDs are passed, and only the highest cost of any individual component route is applied to the aggregate for each ID. If multiple component routes contain the same ID, the highest configured cost is applied to the route. For example, the following two component routes are configured with the cost community attribute using an inbound route policy:

- 10.0.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=100
- 192.168.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=200

If these component routes are aggregated or configured as a multipath, the cost value 200 is advertised, because it has the highest cost.

If one or more component routes do not carry the cost community attribute or the component routes are configured with different IDs, then the default value (2147483647) is advertised for the aggregate or

multipath route. For example, the following three component routes are configured with the cost community attribute using an inbound route policy. However, the component routes are configured with two different IDs.

- 10.0.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=100
- 172.16.0.1
  - POI=IGP
  - cost community ID=2
  - cost number=100
- 192.168.0.1
  - POI=IGP
  - cost community ID=1
  - cost number=200


The single advertised path includes the aggregate cost communities as follows:

{POI=IGP, ID=1, Cost=2147483647} {POI=IGP, ID=2, Cost=2147483647}

## Influencing Route Preference in a Multiexit IGP Network

This figure shows an IGP network with two autonomous system boundary routers (ASBRs) on the edge. Each ASBR has an equal cost path to network 10.8/16.

**Figure 1: Multiexit Point IGP Network**



Both paths are considered to be equal by BGP. If multipath loadsharing is configured, both paths to the routing table are installed and are used to balance the load of traffic. If multipath load balancing is not configured,

the BGP selects the path that was learned first as the best path and installs this path to the routing table. This behavior may not be desirable under some conditions. For example, the path is learned from ISP1 PE2 first, but the link between ISP1 PE2 and ASBR1 is a low-speed link.

The configuration of the cost community attribute can be used to influence the BGP best-path selection process by applying a lower-cost community value to the path learned by ASBR2. For example, the following configuration is applied to ASBR2:

```
RP/0/0/CPU0:router(config)# route-policy ISP2_PE1
RP/0/0/CPU0:router(config-rpl)# set extcommunity cost (1:1)
```

The preceding route policy applies a cost community number of 1 to the 10.8.0.0 route. By default, the path learned from ASBR1 is assigned a cost community number of 2147483647. Because the path learned from ASBR2 has a lower-cost community number, the path is preferred.

## BGP Cost Community Support for EIGRP MPLS VPN PE-CE with Back-door Links


Back-door links in an EIGRP MPLS VPN topology is preferred by BGP if the back-door link is learned first. (A back-door link, or route, is a connection that is configured outside of the VPN between a remote and main site; for example, a WAN leased line that connects a remote site to the corporate network.)

The “prebest path” point of insertion (POI) in the BGP cost community feature supports mixed EIGRP VPN network topologies that contain VPN and back-door links. This POI is applied automatically to EIGRP routes that are redistributed into BGP. The “prebest path” POI carries the EIGRP route type and metric. This POI influences the best-path calculation process by influencing BGP to consider the POI before any other comparison step. No configuration is required. This feature is enabled automatically for EIGRP VPN sites when Cisco IOS XR software is installed on a PE, CE, or back-door router.

For information about configuring EIGRP MPLS VPNs, see the *Cisco IOS XR MPLS Configuration Guide for the Cisco XR 12000 Series Router*.

This figure shows how cost community can be used to support backdoor links in a network.

**Figure 2: Network Showing How Cost Community Can be Used to Support Backdoor Links**



The following sequence of events happens in PE1:

- 1 PE1 learns IPv4 prefix 10.1.1.0/24 from CE1 through EIGRP running a virtual routing and forwarding (VRF) instance. EIGRP selects and installs the best path in the RIB. It also encodes the cost-extended community and adds the information to the RIB.
- 2 The route is redistributed into BGP (assuming that IGP-to-BGP redistribution is configured). BGP also receives the cost-extended community from the route through the redistribution process.

- 3 After BGP has determined the best path for the newly redistributed prefix, the path is advertised to PE peers (PE2).
- 4 PE2 receives the BGP VPNv4 prefix route\_distinguisher:10.1.1.0/24 along with the cost community. It is likely that CE2 advertises the same prefix (because of the back-door link between CE1 and CE2) to PE2 through EIGRP. PE2 BGP would have already learned the CE route through the redistribution process along with the cost community value
- 5 PE2 has two paths within BGP: one with cost community cost1 through multipath BGP (PE1) and another with cost community cost2 through the EIGRP neighbor (CE2).
- 6 PE2 runs the enhanced BGP best-path calculation.
- 7 PE2 installs the best path in the RIB passing the appropriate cost community value.
- 8 PE2 RIB has two paths for 10.1.1.0/24: one with cost community cost2 added by EIGRP and another with the cost community cost1 added by BGP. Because both the route paths have cost community, RIB compares the costs first. The BGP path has the lower cost community, so it is selected and downloaded to the RIB.
- 9 PE2 RIB redistributes the BGP path into EIGRP with VRF. EIGRP runs a diffusing update algorithm (DUAL) because there are two paths, and selects the BGP-redistributed path.
- 10 PE2 EIGRP advertises the path to CE2 making the path the next hop for the prefix to send the traffic over the MPLS network.

## Adding Routes to the Routing Information Base

If a nonsourced path becomes the best path after the best-path calculation, BGP adds the route to the Routing Information Base (RIB) and passes the cost communities along with the other IGP extended communities.

When a route with paths is added to the RIB by a protocol, RIB checks the current best paths for the route and the added paths for cost extended communities. If cost-extended communities are found, the RIB compares the set of cost communities. If the comparison does not result in a tie, the appropriate best path is chosen. If the comparison results in a tie, the RIB proceeds with the remaining steps of the best-path algorithm. If a cost community is not present in either the current best paths or added paths, then the RIB continues with the remaining steps of the best-path algorithm. See [BGP Best Path Algorithm, page 28](#) for information on the BGP best-path algorithm.

## BGP Best Path Algorithm

BGP routers typically receive multiple paths to the same destination. The BGP best-path algorithm determines the best path to install in the IP routing table and to use for forwarding traffic. This section describes the Cisco IOS XR software implementation of BGP best-path algorithm, as specified in Section 9.1 of the Internet Engineering Task Force (IETF) Network Working Group draft-ietf-idr-bgp4-24.txt document.

The BGP best-path algorithm implementation is in three parts:

- Part 1—Compares two paths to determine which is better.
- Part 2—Iterates over all paths and determines which order to compare the paths to select the overall best path.
- Part 3—Determines whether the old and new best paths differ enough so that the new best path should be used.




---

**Note** The order of comparison determined by Part 2 is important because the comparison operation is not transitive; that is, if three paths, A, B, and C exist, such that when A and B are compared, A is better, and when B and C are compared, B is better, it is not necessarily the case that when A and C are compared, A is better. This nontransitivity arises because the multi exit discriminator (MED) is compared only among paths from the same neighboring autonomous system (AS) and not among all paths.

---

## Comparing Pairs of Paths

Perform the following steps to compare two paths and determine the better path:

- 1 If either path is invalid (for example, a path has the maximum possible MED value or it has an unreachable next hop), then the other path is chosen (provided that the path is valid).
- 2 If the paths have unequal pre-bestpath cost communities, the path with the lower pre-bestpath cost community is selected as the best path.




---

**Note** See [BGP Cost Community, page 24](#) for details on how cost communities are compared.

---

- 3 If the paths have unequal weights, the path with the highest weight is chosen.




---

**Note** The weight is entirely local to the router, and can be set with the **weight** command or using a routing policy.

---

- 4 If the paths have unequal local preferences, the path with the higher local preference is chosen.




---

**Note** If a local preference attribute was received with the path or was set by a routing policy, then that value is used in this comparison. Otherwise, the default local preference value of 100 is used. The default value can be changed using the **bgp default local-preference** command.

---

- 5 If one of the paths is a redistributed path, which results from a **redistribute** or **network** command, then it is chosen. Otherwise, if one of the paths is a locally generated aggregate, which results from an **aggregate-address** command, it is chosen.




---

**Note** Step 1 through Step 4 implement the “Path Selection with BGP” of RFC 1268.

---

- 6 If the paths have unequal AS path lengths, the path with the shorter AS path is chosen. This step is skipped if **bgp bestpath as-path ignore** command is configured.




---

**Note** When calculating the length of the AS path, confederation segments are ignored, and AS sets count as 1.

---

**Note**

eiBGP specifies internal and external BGP multipath peers. eiBGP allows simultaneous use of internal and external paths.

- 7 If the paths have different origins, the path with the lower origin is selected. Interior Gateway Protocol (IGP) is considered lower than EGP, which is considered lower than INCOMPLETE.
- 8 If appropriate, the MED of the paths is compared. If they are unequal, the path with the lower MED is chosen.

A number of configuration options exist that affect whether or not this step is performed. In general, the MED is compared if both paths were received from neighbors in the same AS; otherwise the MED comparison is skipped. However, this behavior is modified by certain configuration options, and there are also some corner cases to consider.

If the **bgp bestpath med always** command is configured, then the MED comparison is always performed, regardless of neighbor AS in the paths. Otherwise, MED comparison depends on the AS paths of the two paths being compared, as follows:

- If a path has no AS path or the AS path starts with an AS\_SET, then the path is considered to be internal, and the MED is compared with other internal paths.
- If the AS path starts with an AS\_SEQUENCE, then the neighbor AS is the first AS number in the sequence, and the MED is compared with other paths that have the same neighbor AS.
- If the AS path contains only confederation segments or starts with confederation segments followed by an AS\_SET, then the MED is not compared with any other path unless the **bgp bestpath med confed** command is configured. In that case, the path is considered internal and the MED is compared with other internal paths.
- If the AS path starts with confederation segments followed by an AS\_SEQUENCE, then the neighbor AS is the first AS number in the AS\_SEQUENCE, and the MED is compared with other paths that have the same neighbor AS.

**Note**

If no MED attribute was received with the path, then the MED is considered to be 0 unless the **bgp bestpath med missing-as-worst** command is configured. In that case, if no MED attribute was received, the MED is considered to be the highest possible value.

- 9 If one path is received from an external peer and the other is received from an internal (or confederation) peer, the path from the external peer is chosen.
- 10 If the paths have different IGP metrics to their next hops, the path with the lower IGP metric is chosen.
- 11 If the paths have unequal IP cost communities, the path with the lower IP cost community is selected as the best path.

**Note**

See the [BGP Cost Community](#), page 24 for details on how cost communities are compared.

- 12 If all path parameters in Step 1 through Step 10 are the same, then the router IDs are compared. If the path was received with an originator attribute, then that is used as the router ID to compare; otherwise, the

router ID of the neighbor from which the path was received is used. If the paths have different router IDs, the path with the lower router ID is chosen.

**Note**

Where the originator is used as the router ID, it is possible to have two paths with the same router ID. It is also possible to have two BGP sessions with the same peer router, and therefore receive two paths with the same router ID.

- 13 If the paths have different cluster lengths, the path with the shorter cluster length is selected. If a path was not received with a cluster list attribute, it is considered to have a cluster length of 0.
- 14 Finally, the path received from the neighbor with the lower IP address is chosen. Locally generated paths (for example, redistributed paths) are considered to have a neighbor IP address of 0.

## Order of Comparisons

The second part of the BGP best-path algorithm implementation determines the order in which the paths should be compared. The order of comparison is determined as follows:

- 1 The paths are partitioned into groups such that within each group the MED can be compared among all paths. The same rules as in [Comparing Pairs of Paths, page 29](#) are used to determine whether MED can be compared between any two paths. Normally, this comparison results in one group for each neighbor AS. If the **bgp bestpath med always** command is configured, then there is just one group containing all the paths.
- 2 The best path in each group is determined. Determining the best path is achieved by iterating through all paths in the group and keeping track of the best one seen so far. Each path is compared with the best-so-far, and if it is better, it becomes the new best-so-far and is compared with the next path in the group.
- 3 A set of paths is formed containing the best path selected from each group in Step 2. The overall best path is selected from this set of paths, by iterating through them as in Step 2.

## Best Path Change Suppression

The third part of the implementation is to determine whether the best-path change can be suppressed or not—whether the new best path should be used, or continue using the existing best path. The existing best path can continue to be used if the new one is identical to the point at which the best-path selection algorithm becomes arbitrary (if the router-id is the same). Continuing to use the existing best path can avoid churn in the network.

**Note**

This suppression behavior does not comply with the IETF Networking Working Group draft-ietf-idr-bgp4-24.txt document, but is specified in the IETF Networking Working Group draft-ietf-idr-avoid-transition-00.txt document.

The suppression behavior can be turned off by configuring the **bgp bestpath compare-routerid** command. If this command is configured, the new best path is always preferred to the existing one.

Otherwise, the following steps are used to determine whether the best-path change can be suppressed:

- 1 If the existing best path is no longer valid, the change cannot be suppressed.

- 2 If either the existing or new best paths were received from internal (or confederation) peers or were locally generated (for example, by redistribution), then the change cannot be suppressed. That is, suppression is possible only if both paths were received from external peers.
- 3 If the paths were received from the same peer (the paths would have the same router-id), the change cannot be suppressed. The router ID is calculated using rules in [Comparing Pairs of Paths, page 29](#).
- 4 If the paths have different weights, local preferences, origins, or IGP metrics to their next hops, then the change cannot be suppressed. Note that all these values are calculated using the rules in [Comparing Pairs of Paths, page 29](#).
- 5 If the paths have different-length AS paths and the **bgp bestpath as-path ignore** command is not configured, then the change cannot be suppressed. Again, the AS path length is calculated using the rules in [Comparing Pairs of Paths, page 29](#).
- 6 If the MED of the paths can be compared and the MEDs are different, then the change cannot be suppressed. The decision as to whether the MEDs can be compared is exactly the same as the rules in [Comparing Pairs of Paths, page 29](#), as is the calculation of the MED value.
- 7 If all path parameters in Step 1 through Step 6 do not apply, the change can be suppressed.

## Administrative Distance

An administrative distance is a rating of the trustworthiness of a routing information source. In general, the higher the value, the lower the trust rating. For information on specifying the administrative distance for BGP, see the BGP Commands module of the *Cisco IOS XR Routing Command Reference for the Cisco XR 12000 Series Router*

Normally, a route can be learned through more than one protocol. Administrative distance is used to discriminate between routes learned from more than one protocol. The route with the lowest administrative distance is installed in the IP routing table. By default, BGP uses the administrative distances shown in [Table 1: BGP Default Administrative Distances, page 32](#).

**Table 1: BGP Default Administrative Distances**

Distance	Default Value	Function
External	20	Applied to routes learned from eBGP.
Internal	200	Applied to routes learned from iBGP.
Local	200	Applied to routes originated by the router.



### Note


Distance does not influence the BGP path selection algorithm, but it does influence whether BGP-learned routes are installed in the IP routing table.

In most cases, when a route is learned through eBGP, it is installed in the IP routing table because of its distance (20). Sometimes, however, two ASs have an IGP-learned back-door route and an eBGP-learned



route. Their policy might be to use the IGP-learned path as the preferred path and to use the eBGP-learned path when the IGP path is down. See [Figure 3: Back Door Example](#), page 33.

**Figure 3: Back Door Example**



In [Figure 3: Back Door Example](#), page 33, Routers A and C and Routers B and C are running eBGP. Routers A and B are running an IGP (such as Routing Information Protocol [RIP], Interior Gateway Routing Protocol [IGRP], Enhanced IGRP, or Open Shortest Path First [OSPF]). The default distances for RIP, IGRP, Enhanced IGRP, and OSPF are 120, 100, 90, and 110, respectively. All these distances are higher than the default distance of eBGP, which is 20. Usually, the route with the lowest distance is preferred.

Router A receives updates about 160.10.0.0 from two routing protocols: eBGP and IGP. Because the default distance for eBGP is lower than the default distance of the IGP, Router A chooses the eBGP-learned route from Router C. If you want Router A to learn about 160.10.0.0 from Router B (IGP), establish a BGP back door. See .

In the following example, a network back-door is configured:

```
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)# network 160.10.0.0/16 backdoor
```

Router A treats the eBGP-learned route as local and installs it in the IP routing table with a distance of 200. The network is also learned through Enhanced IGRP (with a distance of 90), so the Enhanced IGRP route is successfully installed in the IP routing table and is used to forward traffic. If the Enhanced IGRP-learned route goes down, the eBGP-learned route is installed in the IP routing table and is used to forward traffic.

Although BGP treats network 160.10.0.0 as a local entry, it does not advertise network 160.10.0.0 as it normally would advertise a local entry.

## Multiprotocol BGP

Multiprotocol BGP is an enhanced BGP that carries routing information for multiple network layer protocols and IP multicast routes. BGP carries two sets of routes, one set for unicast routing and one set for multicast routing. The routes associated with multicast routing are used by the Protocol Independent Multicast (PIM) feature to build data distribution trees.

Multiprotocol BGP is useful when you want a link dedicated to multicast traffic, perhaps to limit which resources are used for which traffic. Multiprotocol BGP allows you to have a unicast routing topology different from a multicast routing topology providing more control over your network and resources.

In BGP, the only way to perform interdomain multicast routing was to use the BGP infrastructure that was in place for unicast routing. Perhaps you want all multicast traffic exchanged at one network access point (NAP). If those routers were not multicast capable, or there were differing policies for which you wanted multicast traffic to flow, multicast routing could not be supported without multiprotocol BGP.

**Note**

It is possible to configure BGP peers that exchange both unicast and multicast network layer reachability information (NLRI), but you cannot connect multiprotocol BGP clouds with a BGP cloud. That is, you cannot redistribute multiprotocol BGP routes into BGP.

Figure 4: [Noncongruent Unicast and Multicast Routes](#), page 34 illustrates simple unicast and multicast topologies that are incongruent, and therefore are not possible without multiprotocol BGP.

Autonomous systems 100, 200, and 300 are each connected to two NAPs that are FDDI rings. One is used for unicast peering (and therefore the exchange of unicast traffic). The Multicast Friendly Interconnect (MFI) ring is used for multicast peering (and therefore the exchange of multicast traffic). Each router is unicast and multicast capable.

**Figure 4: Noncongruent Unicast and Multicast Routes**





Figure 5: [Multicast BGP Environment](#), page 35 is a topology of unicast-only routers and multicast-only routers. The two routers on the left are unicast-only routers (that is, they do not support or are not configured to perform multicast routing). The two routers on the right are multicast-only routers. Routers A and B support both unicast and multicast routing. The unicast-only and multicast-only routers are connected to a single NAP.

In Figure 5: [Multicast BGP Environment](#), page 35, only unicast traffic can travel from Router A to the unicast routers to Router B and back. Multicast traffic could not flow on that path, so another routing table is required. Multicast traffic uses the path from Router A to the multicast routers to Router B and back.

Figure 5: [Multicast BGP Environment](#), page 35 illustrates a multiprotocol BGP environment with a separate unicast route and multicast route from Router A to Router B. Multiprotocol BGP allows these routes to be incongruent. Both of the autonomous systems must be configured for internal multiprotocol BGP (IMBGP) in the figure.

A multicast routing protocol, such as PIM, uses the multicast BGP database to perform Reverse Path Forwarding (RPF) lookups for multicast-capable sources. Thus, packets can be sent and accepted on the multicast topology but not on the unicast topology.

**Figure 5: Multicast BGP Environment**



## Route Dampening

Route dampening is a BGP feature that minimizes the propagation of flapping routes across an internetwork. A route is considered to be flapping when it is repeatedly available, then unavailable, then available, then unavailable, and so on.

For example, consider a network with three BGP autonomous systems: autonomous system 1, autonomous system 2, and autonomous system 3. Suppose the route to network A in autonomous system 1 flaps (it becomes unavailable). Under circumstances without route dampening, the eBGP neighbor of autonomous system 1 to autonomous system 2 sends a withdraw message to autonomous system 2. The border router in autonomous system 2, in turn, propagates the withdrawal message to autonomous system 3. When the route to network A reappears, autonomous system 1 sends an advertisement message to autonomous system 2, which sends it to autonomous system 3. If the route to network A repeatedly becomes unavailable, then available, many withdrawal and advertisement messages are sent. Route flapping is a problem in an internetwork connected to the Internet, because a route flap in the Internet backbone usually involves many routes.

## Minimizing Flapping

The route dampening feature minimizes the flapping problem as follows. Suppose again that the route to network A flaps. The router in autonomous system 2 (in which route dampening is enabled) assigns network A a penalty of 1000 and moves it to history state. The router in autonomous system 2 continues to advertise the status of the route to neighbors. The penalties are cumulative. When the route flaps so often that the penalty exceeds a configurable suppression limit, the router stops advertising the route to network A, regardless of how many times it flaps. Thus, the route is dampened.

The penalty placed on network A is decayed until the reuse limit is reached, upon which the route is once again advertised. At half of the reuse limit, the dampening information for the route to network A is removed.

**Note**

---

No penalty is applied to a BGP peer reset when route dampening is enabled, even though the reset withdraws the route.

---

## BGP Routing Domain Confederation

One way to reduce the iBGP mesh is to divide an autonomous system into multiple subautonomous systems and group them into a single confederation. To the outside world, the confederation looks like a single autonomous system. Each autonomous system is fully meshed within itself and has a few connections to other autonomous systems in the same confederation. Although the peers in different autonomous systems have eBGP sessions, they exchange routing information as if they were iBGP peers. Specifically, the next hop, MED, and local preference information is preserved. This feature allows you to retain a single IGP for all of the autonomous systems.

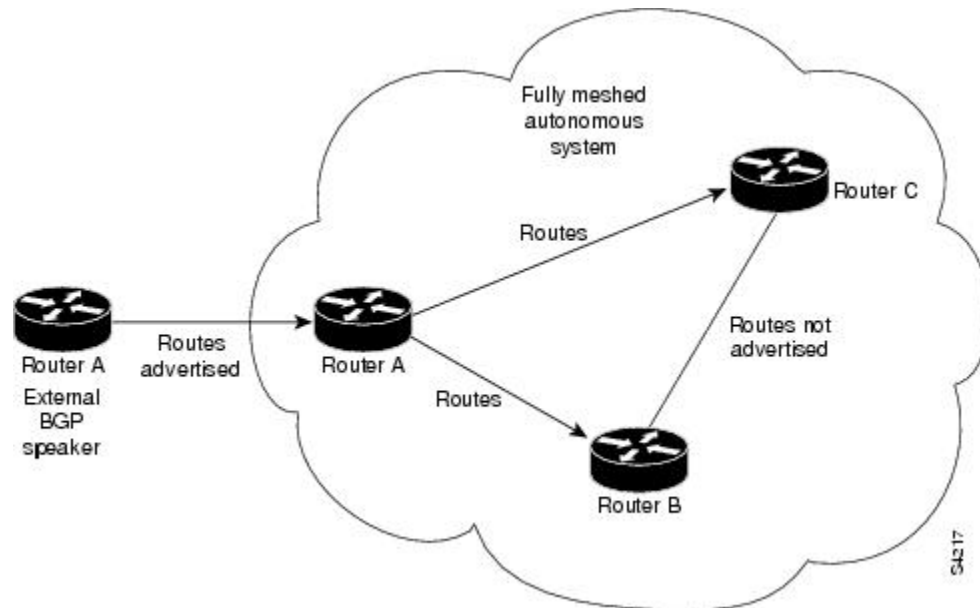
## BGP Route Reflectors

BGP requires that all iBGP speakers be fully meshed. However, this requirement does not scale well when there are many iBGP speakers. Instead of configuring a confederation, you can reduce the iBGP mesh by using a route reflector configuration.

[Figure 6: Three Fully Meshed iBGP Speakers, page 37](#) illustrates a simple iBGP configuration with three iBGP speakers (routers A, B, and C). Without route reflectors, when Router A receives a route from an external neighbor, it must advertise it to both routers B and C. Routers B and C do not readvertise the iBGP learned

route to other iBGP speakers because the routers do not pass on routes learned from internal neighbors to other internal neighbors, thus preventing a routing information loop.


**Figure 6: Three Fully Meshed iBGP Speakers**



With route reflectors, all iBGP speakers need not be fully meshed because there is a method to pass learned routes to neighbors. In this model, an iBGP peer is configured to be a route reflector responsible for passing iBGP learned routes to a set of iBGP neighbors. In [Figure 7: Simple BGP Model with a Route Reflector, page 38](#), Router B is configured as a route reflector. When the route reflector receives routes advertised from

Router A, it advertises them to Router C, and vice versa. This scheme eliminates the need for the iBGP session between routers A and C.

**Figure 7: Simple BGP Model with a Route Reflector**



The internal peers of the route reflector are divided into two groups: client peers and all other routers in the autonomous system (nonclient peers). A route reflector reflects routes between these two groups. The route reflector and its client peers form a *cluster*. The nonclient peers must be fully meshed with each other, but the

client peers need not be fully meshed. The clients in the cluster do not communicate with iBGP speakers outside their cluster.

**Figure 8: More Complex BGP Route Reflector Model**

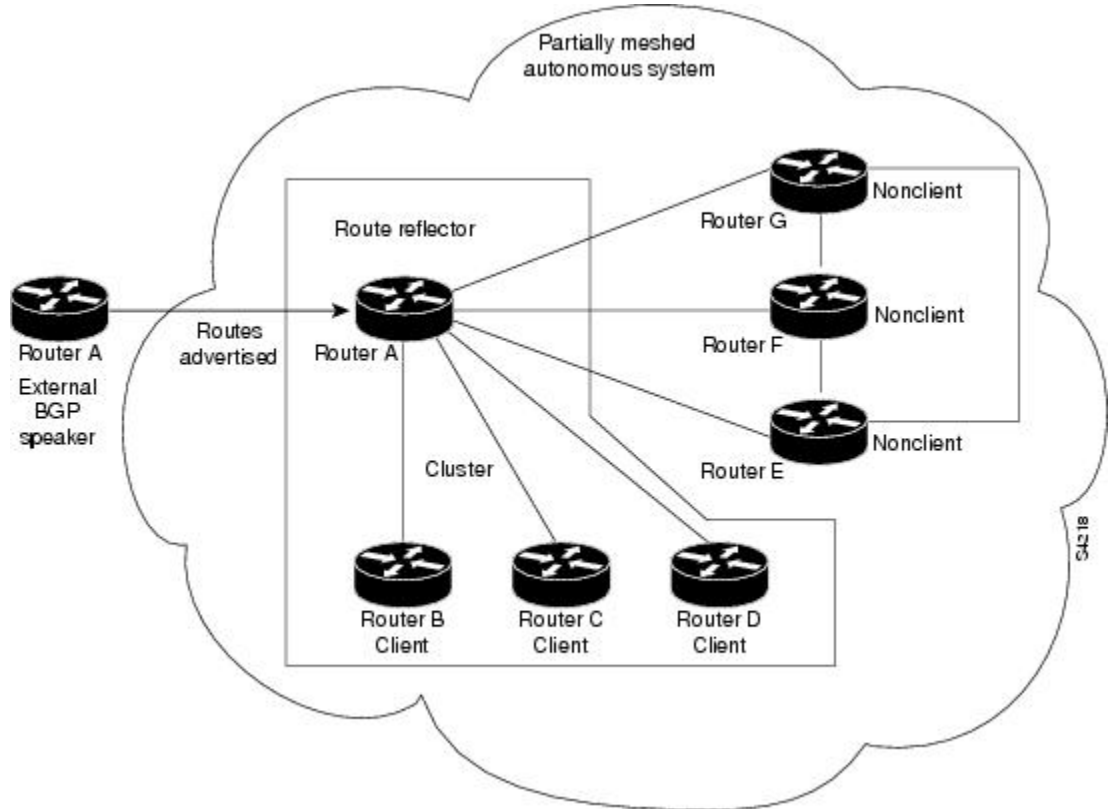


Figure 8: More Complex BGP Route Reflector Model, page 39 illustrates a more complex route reflector scheme. Router A is the route reflector in a cluster with routers B, C, and D. Routers E, F, and G are fully meshed, nonclient routers.

When the route reflector receives an advertised route, depending on the neighbor, it takes the following actions:

- A route from an external BGP speaker is advertised to all clients and nonclient peers.
- A route from a nonclient peer is advertised to all clients.
- A route from a client is advertised to all clients and nonclient peers. Hence, the clients need not be fully meshed.

Along with route reflector-aware BGP speakers, it is possible to have BGP speakers that do not understand the concept of route reflectors. They can be members of either client or nonclient groups, allowing an easy and gradual migration from the old BGP model to the route reflector model. Initially, you could create a single cluster with a route reflector and a few clients. All other iBGP speakers could be nonclient peers to the route reflector and then more clusters could be created gradually.

An autonomous system can have multiple route reflectors. A route reflector treats other route reflectors just like other iBGP speakers. A route reflector can be configured to have other route reflectors in a client group or nonclient group. In a simple configuration, the backbone could be divided into many clusters. Each route

reflector would be configured with other route reflectors as nonclient peers (thus, all route reflectors are fully meshed). The clients are configured to maintain iBGP sessions with only the route reflector in their cluster.

Usually, a cluster of clients has a single route reflector. In that case, the cluster is identified by the router ID of the route reflector. To increase redundancy and avoid a single point of failure, a cluster might have more than one route reflector. In this case, all route reflectors in the cluster must be configured with the cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster. All route reflectors serving a cluster should be fully meshed and all of them should have identical sets of client and nonclient peers.

By default, the clients of a route reflector 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, the route reflector need not reflect routes to clients.

As the iBGP learned routes are reflected, routing information may loop. The route reflector model has the following mechanisms to avoid routing loops:

- Originator ID is an optional, nontransitive BGP attribute. It is a 4-byte attributed created by a route reflector. The attribute carries the router ID of the originator of the route in the local autonomous system. Therefore, if a misconfiguration causes routing information to come back to the originator, the information is ignored.
- Cluster-list is an optional, nontransitive BGP attribute. It is a sequence of cluster IDs that the route has passed. When a route reflector reflects a route from its clients to nonclient peers, and vice versa, it appends the local cluster ID to the cluster-list. If the cluster-list is empty, a new cluster-list is created. Using this attribute, a route reflector can identify if routing information is looped back to the same cluster due to misconfiguration. If the local cluster ID is found in the cluster-list, the advertisement is ignored.

## Default Address Family for show Commands

Most of the **show** commands provide address family (AFI) and subaddress family (SAFI) arguments (see RFC 1700 and RFC 2858 for information on AFI and SAFI). The Cisco IOS XR software parser provides the ability to set the **afi** and **safi** so that it is not necessary to specify them while running a **show** command. The parser commands are:

- **set default-afi { ipv4 | ipv6 | all }**
- **set default-safi { unicast | multicast | all }**

The parser automatically sets the default **afi** value to **ipv4** and default **safi** value to **unicast**. It is necessary to use only the parser commands to change the default **afi** value from **ipv4** or default **safi** value from **unicast**. Any **afi** or **safi** keyword specified in a **show** command overrides the values set using the parser commands. Use the following **show default-afi-safi-vrf** command to check the currently set value of the **afi** and **safi**.

## Distributed BGP

Distributed BGP splits BGP functionality into three process types:

- BGP process manager—Responsible for verifying configuration changes and for calculating and publishing the distribution of neighbors among BGP speaker processes.

There is a single instance of this process.



- bRIB process—Responsible for performing the best-path calculation of routes (receives partial best paths from the speaker). The best route is installed into the bRIB and is advertised back to all speakers. See the [BGP Best Path Algorithm, page 28](#) for information on best-path calculation. The bRIB process is also responsible for installing routes in the RIB, and for handling routes redistributed from the RIB. To accommodate route leaking from one RIB to another, bRIB may register for redistribution from multiple RIB routes into a single route in the bRIB process.

There is a single instance of this process for each address family.

- BGP speaker process—Responsible for handling all BGP connections to peers. The speaker stores received paths in the RIB and performs a partial best-path calculation, advertising the partial best paths to the bRIB (limited best-path calculation). Speakers perform a limited best-path calculation because to compare Multi Exit Discriminators (MEDs), paths need to be compared from the same AS but may not be received on the same speaker. Because BGP speakers do not have access to the entire BGP local RIB, BGP speakers can perform only a limited best-path calculation. (These are Step 1 through Step 7 in the [BGP Best Path Algorithm, page 28](#).) Only the best paths are advertised to the bRIB to reduce speaker/bRIB interprocess communications (IPC) and to reduce the number of paths to be processed in the bRIB. BGP speakers can only mark a path as active only after learning the result of the full best-path calculation from the bRIB. Neighbor import and export policies are imposed by the speaker.

If the **bgp bestpath med always** command is enabled, complete best-path calculation happens inside speaker process. When the **bgp bestpath med always** command is not enabled, speakers calculate partial best paths only (performs the best-path steps up to the MED comparison) and send them to bRIB. bRIB calculates the final best path (performs all the steps in the best-path calculation). When the **bgp bestpath med always** command is enabled, 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.

There are multiple instances of this process in which each instance is responsible for a subset of BGP peer connections.

Up to a total 15 speakers for all address families and one bRIB for each address family (IPv4, IPv6, and VPNv4) are supported.

Distributed BGP is used to reduce the impact that a fault in one address family has on another address family. For example, you can have one speaker with only IPv6 neighbors (peering to IPv6 addresses) and a separate speaker with only IPv4 neighbors (peering to IPv4 addresses), and yet another speaker with only VPNv4 provider edge (PE) or customer edge (CE) neighbors (peering to IPv4 addresses distinct from the non-VPN neighbors). In this scenario, there is no overlap in processes (bgp, bRIB, and RIB) between IPv4, IPv6, and VPNv4. Therefore, a bgp, bRIB, or RIB process crash affects only one address family. Distributed BGP also allows more CPU capacity for receiving, computing, and sending BGP routing updates. When in distributed BGP mode, you can control the number of distributed speakers that are enabled, as well as which neighbors are assigned to each speaker. If no distributed speakers are enabled, BGP operates in standalone mode. If at least one distributed speaker is enabled, BGP operates in distributed mode.

## MPLS VPN Carrier Supporting Carrier

Carrier supporting carrier (CSC) is a term used to describe a situation in which one service provider allows another service provider to use a segment of its backbone network. The service provider that provides the segment of the backbone network to the other provider is called the *backbone carrier*. The service provider that uses the segment of the backbone network is called the *customer carrier*.

A backbone carrier offers Border Gateway Protocol and Multiprotocol Label Switching (BGP/MPLS) VPN services. The customer carrier can be either:

- An Internet service provider (ISP) (By definition, an ISP does not provide VPN service.)
- A BGP/MPLS VPN service provider

You can configure a CSC network to enable BGP to transport routes and MPLS labels between the backbone carrier provider edge (PE) routers and the customer carrier customer edge (CE) routers using multiple paths. The benefits of using BGP to distribute IPv4 routes and MPLS label routes are:

- BGP takes the place of an Interior Gateway Protocol (IGP) and Label Distribution Protocol (LDP) in a VPN routing and forwarding (VRF) table. You can use BGP to distribute routes and MPLS labels. Using a single protocol instead of two simplifies the configuration and troubleshooting.
- BGP is the preferred routing protocol for connecting two ISPs, mainly because of its routing policies and ability to scale. ISPs commonly use BGP between two providers. This feature enables those ISPs to use BGP.

For detailed information on configuring MPLS VPN CSC with BGP, see the *Implementing MPLS Layer 3 VPNs on Cisco IOS XR Software* module of the *Cisco IOS XR MPLS Configuration Guide for the Cisco XR 12000 Series Router*.

## BGP Keychains

BGP keychains enable keychain authentication between two BGP peers. The BGP endpoints must both comply with draft-bonica-tcp-auth-05.txt and a keychain on one endpoint and a password on the other endpoint does not work.

See the *Cisco IOS XR System Security Configuration Guide for the Cisco XR 12000 Series Router* for information on keychain management.

BGP is able to use the keychain to implement hitless key rollover for authentication. Key rollover specification is time based, and in the event of clock skew between the peers, the rollover process is impacted. The configurable tolerance specification allows for the accept window to be extended (before and after) by that margin. This accept window facilitates a hitless key rollover for applications (for example, routing and management protocols).

The key rollover does not impact the BGP session, unless there is a keychain configuration mismatch at the endpoints resulting in no common keys for the session traffic (send or accept).

## IPv6/IPv6 VPN Provider Edge Transport over MPLS

IPv6 Provider Edge (6PE) and IPv6 VPN Provider Edge (6VPE) leverages the existing Multiprotocol Label Switching (MPLS) IPv4 core infrastructure for IPv6 transport. 6PE and 6VPE enables IPv6 sites to communicate with each other over an MPLS IPv4 core network using MPLS label switched paths (LSPs). This feature relies on multiprotocol Border Gateway Protocol (BGP) extensions in the IPv4 network configuration on the provider edge (PE) router, to exchange IPv6 reachability information in addition to an MPLS label for each IPv6 address prefix to be advertised. Edge routers are configured to be dual stack running both IPv4 and IPv6, and use the IPv4-mapped IPv6 address for IPv6 prefix reachability exchange.

For detailed information on configuring 6PE and 6VPE over MPLS, see *Cisco IOS XR MPLS Configuration Guide for the Cisco XR 12000 Series Router*.

## IPv6 Provider Edge Multipath

Internal and external BGP multipath for IPv6 allows the IPv6 router to load balance between several paths (for example, same neighboring autonomous system [AS] or sub-AS, or the same metric) to reach its destination. The 6PE multipath feature uses multiprotocol internal BGP (MP-iBGP) to distribute IPv6 routes over the MPLS IPv4 core network and to attach an MPLS label to each route.

When MP-iBGP multipath is enabled on the 6PE router, all labeled paths are installed in the forwarding table with MPLS information (label stack) when MPLS information is available. This functionality enables 6PE to perform load balancing.

## VPNv4/VPNv6 over the IP Core Using L2TPv3 Tunnels

The Layer 2 Tunnel Protocol version 3 (L2TPv3) feature defines the L2TP protocol for tunneling Layer 2 traffic over an IP core network using Layer 2 VPNs. Benefits of this feature include:


- Simplifies deployment of VPNs
- Does not require Multiprotocol Label Switching (MPLS)
- Supports Layer 2 tunneling over IP for any traffic
- Supports data encapsulation directly over IP (IP protocol number 115), not using User Datagram Protocol (UDP)
- Supports point-to-point sessions, not point-to-multipoint or multipoint-to-point sessions
- Supports sessions between the same Layer 2 protocols, for example Frame Relay to Frame Relay or ATM to ATM

When an RFC 4364-based IP VPN service is deployed (see RFC 4364), VPN traffic is typically transported across the core network between service provider edge (PE) routers using MPLS label switched paths (LSPs). Native IP L3VPNs eliminate the need for MPLS between the participating core routers by relying on scalable tunnel encapsulation over IP. These tunnels can be used instead of, or with, MPLS to transport VPN traffic between participating edge routers.

A native IP L3VPN allows service providers to use an IP backbone to provide VPN services. BGP is used to distribute VPN routing information across the provider backbone.

This figure shows edge routers participating in switching IPv4 and IPv6 traffic over a tunnel using IP as the transport.

Figure 9: IP L3VPN with L2TPv3



## BGP Multicast VPN

The BGP Multicast VPN feature uses the IPv4 multicast distribution tree (MDT) subaddress family identifier (SAFI) in Border Gateway Protocol (BGP).

Multicast VPN (MVPN) extends the VPN architecture to provide multicast services over a shared service provider backbone using native multicast technology. This is achieved using virtual connections between provider edge (PE) routers in each VPN and using native multicast forwarding inside the provider network. An MDT may span across multiple customer sites and the provider network, allowing traffic to flow freely from one source to multiple receivers.

MVPN is supported on VPN networks based on MPLS and on networks based on IP Layer 2 Tunnel Protocol version 3 (L2TPv3).

PE routers are the only routers that must be MVPN-aware and that must be able to signal to remote PEs information regarding the MVPN. Therefore, all PE routers must have a BGP relationship with each other—either directly or using a route reflector (RR).

Generally the source address of the default MDT is the same address used to source the internal BGP (iBGP) sessions with the remote PE routers that belong to the same VPN and multicast VPN routing and forwarding (MVRF) instance. When Protocol Independent Multicast–Source Specific Multicast (PIM–SSM) is used for transport inside the provider core, it is through the BGP relationship that the PEs indicate that they are MVPN-capable and provide for source discovery. This capability is indicated using the updated BGP message.



### Note

The source address can also be configured uniquely per VRF instance under multicast-routing configuration. See *Cisco IOS XR Multicast Configuration Guide for the Cisco XR 12000 Series Router*.

When a PE receives a BGP update, which includes the rendezvous point (RP) and the group information, it joins the root of that tree, thereby joining the MDT.

Figure 10: Multiprotocol iBGP Updates for MVPN, page 45 shows Multiprotocol iBGP updates for MVPN. On PE1, PE2 is configured as its iBGP peer. This BGP peer configuration within a VRF triggers the MP-iBGP updates that send PE1 local VPN routes to PE2. BGP process on PE2 receives the VPN updates and installs VPN routes in the Routing Information Base (RIB) VRF table. When PIM looks up a VRF source or rendezvous point address that is reachable through the provider core, it receives an MP-iBGP route from the RIB.

When an MVPN-specific default MDT group is configured on PE1, PIM creates a virtual MDT tunnel interface with the tunnel source address the same as the BGP local peering address. This MDT interface is used by PIM to send VPN packets to the provider network and to receive VPN packets from the provider network. PIM also exchanges control messages over this MDT interface.

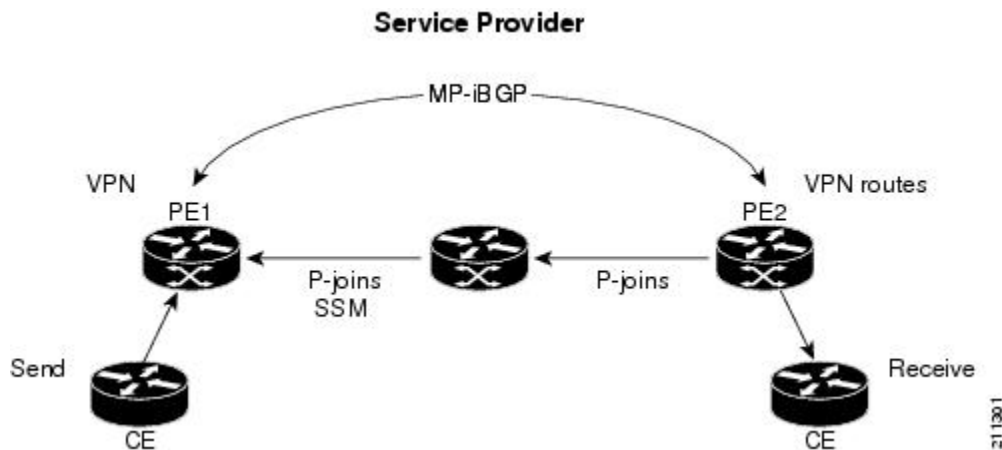
Each time a default MDT group is configured for a specific VRF, BGP builds an MDT SAFI update, with network layer reachability information (NLRI) containing the local PE BGP peering address and the newly configured MDT group address (The NLRI format is 8-byte-RD:IPv4-address followed by the MDT group address). This update is sent to all the BGP peers including PE2. The BGP process on PE2 receives this MDT update and notifies PIM. If the group is a PIM-SSM group, PIM on PE2 begins sending SSM joins to the BGP peering address on PE1 to establish an SSM tree in the core. This SSM tree is used to carry PIM control traffic and multicast data traffic in the corresponding VRF.

In summary, PIM requires the following from BGP:

- A new BGP MDT SAFI, which carries the VRF RD and BGP local peering address and default MDT group in its NLRI.
- A notification mechanism from BGP to PIM about the availability of the MDT SAFI update.
- A notification mechanism from PIM to BGP about the default MDT group address and source address.

See Internet Engineering Task Force (IETF) draft-nalawade-idr-mdt-safi-03 for detailed information on MDT SAFI.

Figure 10: Multiprotocol iBGP Updates for MVPN



## BGP Nonstop Routing

The Border Gateway Protocol (BGP) Nonstop Routing (NSR) with Stateful Switchover (SSO) feature enables all bgp peerings to maintain the BGP state and ensure continuous packet forwarding during events that could

interrupt service. Under NSR, events that might potentially interrupt service are not visible to peer routers. Protocol sessions are not interrupted and routing states are maintained across process restarts and switchovers.

BGP NSR provides nonstop routing during the following events:

- Route processor switchover
- Process restart of BGP or TCP
- In-Service System Upgrade (ISSU)
- Minimum Disruption Restart (MDR)

During route processor switchover and In-Service System Upgrade (ISSU), NSR is achieved by stateful switchover (SSO) of both TCP and BGP.

NSR does not force any software upgrades on other routers in the network, and peer routers are not required to support NSR.

When a route processor switchover occurs due to a fault, the TCP connections and the BGP sessions are migrated transparently to the standby route processor, and the standby route processor becomes active. The existing protocol state is maintained on the standby route processor when it becomes active, and the protocol state does not need to be refreshed by peers.

Events such as soft reconfiguration and policy modifications can trigger the BGP internal state to change. To ensure state consistency between active and standby BGP processes during such events, the concept of post-it is introduced that act as synchronization points.

BGP NSR provides the following features:

- NSR-related alarms and notifications
- Configured and operational NSR states are tracked separately
- NSR statistics collection
- NSR statistics display using **show** commands
- XML schemas support
- Auditing mechanisms to verify state synchronization between active and standby instances
- CLI commands to enable and disable NSR

NSR can be provisioned on a multishelf router, especially if distributed BGP is configured. The following guidelines should be observed when provisioning NSR on a multishelf router:

- When provisioning NSR for line cards installed on a single rack, provision the active and standby applications on the distributed route processor (DRP) of that rack. If a rack failure occurs, sessions are dropped, because all line cards go down.
- When provisioning NSR for line cards installed on different racks, use one of the following three options:
  - Provision the active and standby applications on a distributed route processor (DRP) redundant pair, where there is a separate route processor in each rack. This configuration uses up two revenue-producing line-card slots on each rack, but is the most secure configuration.
  - Provision the active and standby application instances using distributed BGP so that the routing sessions on one rack are serviced by a speaker on that rack. The speaker's standby instance is on a distributed route processor (DRP) on the same rack. If a rack failure occurs, it affects only the sessions on that rack and does not result in NSR loss.

- Provision the active and standby applications on a distributed route processor (DRP) pair that spans two racks. In this configuration, the active/standby role of the line cards is not dependent on the active/standby role of the DRPs. This is called *flexible process redundancy* and provides for rack loss and efficient use of LC slots. Use of distributed BGP is not required with this solution.

**Note**

Sessions on line cards in a lost rack are not protected with any of the above options, because there is no line-card redundancy. These options ensure only that sessions on other racks are not affected by a lost rack. However, lost sessions from a lost rack may cause some traffic loss on other racks, because destinations learned through those lost sessions may no longer have alternate routes. Also, rack loss may cause the CPUs on route processors of active racks to slow as they attempt to define new paths for some routes.

## BGP Best-External Path

The Border Gateway Protocol (BGP) best-external path functionality supports advertisement of the best-external path to the iBGP and Route Reflector peers when a locally selected bestpath is from an internal peer.

BGP selects one best path and one backup path to every destination. By default, selects one best path . Additionally, BGP selects another bestpath from among the remaining external paths for a prefix. Only a single path is chosen as the best-external path and is sent to other PEs as the backup path.

BGP calculates the best-external path only when the best path is an iBGP path. If the best path is an eBGP path, then best-external path calculation is not required.

The procedure to determine the best-external path is as follows:

- 1 Determine the best path from the entire set of paths available for a prefix.
- 2 Eliminate the current best path.
- 3 Eliminate all the internal paths for the prefix.
- 4 From the remaining paths, eliminate all the paths that have the same next hop as that of the current best path.
- 5 Rerun the best path algorithm on the remaining set of paths to determine the best-external path.

BGP considers the external and confederations BGP paths for a prefix to calculate the best-external path.

BGP advertises the best path and the best-external path as follows:

- On the primary PE—advertises the best path for a prefix to both its internal and external peers
- On the backup PE—advertises the best path selected for a prefix to the external peers and advertises the best-external path selected for that prefix to the internal peers

The **advertise best-external** command enables the advertisement of the best-external path in global address family and VRF address family configuration modes.

## BGP Prefix Independent Convergence Unipath Primary/Backup

The Border Gateway Protocol Prefix Independent Convergence Unipath (BGP PIC Unipath) primary/backup feature provides the capability to install a backup path into the forwarding table. Installing the backup path provides prefix independent convergence in the event of a primary PE–CE link failure.

The primary/backup path provides a mechanism for BGP to determine a backup best path. The backup best path acts as a backup to the overall best path, which is the primary best path. BGP programs both the paths into the Forwarding Information Base (FIB).

The procedure to determine the backup best path is as follows:

- 1 Determine the best path from the entire set of paths available for a prefix.
- 2 Eliminate the current best path.
- 3 Eliminate all the paths that have the same next hop as that of the current best path.
- 4 Rerun the best path algorithm on the remaining set of paths to determine the backup best path.

The PE-CE local convergence is in the order of four to five seconds for 10000 prefixes. Installing a backup path on the linecards, so that the Forwarding Information Base (FIB) can immediately switch to an alternate path, in the event of a primary PE-CE link failure reduces the convergence time.

In the case of primary PE-CE link failure, the FIB starts forwarding the received traffic towards the backup PE. FIB will continue forwarding the received traffic towards the backup PE for the duration of the network convergence. Since the approach of using a backup path is independent to the prefixes, Prefix Independent Convergence Unipath functionality provides a prefix independent sub second convergence.

The **additional-paths install backup** command installs the backup path in the Forwarding Information Base (FIB) to enable primary backup path.

## BGP Local Label Retention

When a primary PE-CE link fails, BGP withdraws the route corresponding to the primary path along with its local label and programs the backup path in the Routing Information Base (RIB) and the Forwarding Information Base (FIB), by default.

However, until all the internal peers of the primary PE reconverge to use the backup path as the new bestpath, the traffic continues to be forwarded to the primary PE with the local label that was allocated for the primary path. Hence the previously allocated local label for the primary path must be retained on the primary PE for some configurable time after the reconvergence. BGP Local Label Retention feature enables the retention of the local label for a specified period. If no time is specified, the local label is retained for a default value of five minutes.

The **retain local-label** command enables the retention of the local label until the network is converged.

## Command Line Interface (CLI) Consistency for BGP Commands

From Cisco IOS XR Release 3.9.0 onwards, the Border Gateway Protocol (BGP) commands use **disable** keyword to disable a feature. The keyword **inheritance-disable** disables the inheritance of the feature properties from the parent level.



# How to Implement BGP on Cisco IOS XR Software

## Enabling BGP Routing

Perform this task to enable BGP routing and establish a BGP routing process. Configuring BGP neighbors is included as part of enabling BGP routing.

**Note**

---

At least one neighbor and at least one address family must be configured to enable BGP routing. At least one neighbor with both a remote AS and an address family must be configured globally using the **address family** and **remote as** commands.

---

**Before You Begin**

BGP must be able to obtain a router identifier (for example, a configured loopback address). At least, one address family must be configured in the BGP router configuration and the same address family must also be configured under the neighbor.

**Note**

---

If the neighbor is configured as an external BGP (eBGP) peer, you must configure an inbound and outbound route policy on the neighbor using the **route-policy** command.

---

**SUMMARY STEPS**

1. **configure**
2. **route-policy** *route-policy-name*
3. **end-policy**
4. Do one of the following:
  - **end**
  - **commit**
5. **configure**
6. **router bgp** *as-number*
7. **bgp router-id** *ip-address*
8. **address-family** { **ipv4** | **ipv6** } **unicast**
9. **exit**
10. **neighbor** *ip-address*
11. **remote-as** *as-number*
12. **address-family** { **ipv4** | **ipv6** } **unicast**
13. **route-policy** *route-policy-name* { **in** | **out** }
14. Do one of the following:
  - **end**
  - **commit**

**DETAILED STEPS**

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>route-policy</b> <i>route-policy-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# route-policy drop-as-1234 RP/0/0/CPU0:router(config-rpl)# if as-path passes-through '1234' then RP/0/0/CPU0:router(config-rpl)# apply check-communities RP/0/0/CPU0:router(config-rpl)# else RP/0/0/CPU0:router(config-rpl)# pass RP/0/0/CPU0:router(config-rpl)# endif	(Optional) Creates a route policy and enters route policy configuration mode, where you can define the route policy.

	Command or Action	Purpose
<b>Step 3</b>	<b>end-policy</b>  <b>Example:</b> RP/0/0/CPU0:router(config-rpl)# end-policy	(Optional) Ends the definition of a route policy and exits route policy configuration mode.
<b>Step 4</b>	Do one of the following: <ul style="list-style-type: none"> <li>• end</li> <li>• commit</li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config)# end or RP/0/0/CPU0:router(config)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
<b>Step 5</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 6</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the BGP AS number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 7</b>	<b>bgp router-id <i>ip-address</i></b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp router-id 192.168.70.24	Configures the local router with a specified router ID.

	Command or Action	Purpose
<b>Step 8</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either the IPv4 or IPv6 address family and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 9</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
<b>Step 10</b>	<b>neighbor ip-address</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 11</b>	<b>remote-as as-number</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 12</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either the IPv4 or IPv6 address family and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 13</b>	<b>route-policy route-policy-name { in   out }</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# route-policy drop-as-1234 in	(Optional) Applies the specified policy to inbound IPv4 unicast routes.
<b>Step 14</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# end or RP/0/0/CPU0:router(config-bgp-nbr-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre> Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]: </pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring a Routing Domain Confederation for BGP

Perform this task to configure the routing domain confederation for BGP. This includes specifying a confederation identifier and autonomous systems that belong to the confederation.

Configuring a routing domain confederation reduces the internal BGP (iBGP) mesh by dividing an autonomous system into multiple autonomous systems and grouping 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. 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.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp confederation identifier** *as-number*
4. **bgp confederation peers** *as-number*
5. Do one of the following:
  - **end**
  - **commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b>  RP/0/0/CPU0:router# <code>configure</code>	Enters global configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# router bgp 120</pre>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<p><b>bgp confederation identifier</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# bgp confederation identifier 5</pre>	Specifies a BGP confederation identifier.
<b>Step 4</b>	<p><b>bgp confederation peers</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# bgp confederation peers 1091 RP/0/0/CPU0:router(config-bgp)# bgp confederation peers 1092 RP/0/0/CPU0:router(config-bgp)# bgp confederation peers 1093 RP/0/0/CPU0:router(config-bgp)# bgp confederation peers 1094 RP/0/0/CPU0:router(config-bgp)# bgp confederation peers 1095 RP/0/0/CPU0:router(config-bgp)# bgp confederation peers 1096</pre>	Specifies that the BGP autonomous systems belong to a specified BGP confederation identifier. You can associate multiple AS numbers to the same confederation identifier, as shown in the example.
<b>Step 5</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# end or RP/0/0/CPU0:router(config-bgp)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

Command or Action	Purpose
-------------------	---------

## Resetting an eBGP Session Immediately Upon Link Failure

By default, if a link goes down, all BGP sessions of any directly adjacent external peers are immediately reset. Use the **bgp fast-external-fallover disable** command to disable automatic resetting. Turn the automatic reset back on using the **no bgp fast-external-fallover disable** command.

eBGP sessions flap when the node reaches 3500 eBGP sessions with BGP timer values set as 10 and 30. To support more than 3500 eBGP sessions, increase the packet rate by using the **lpts pifib hardware police location location-id** command. Following is a sample configuration to increase the eBGP sessions:

```
RP/0/0/CPU0:router#configure
RP/0/0/CPU0:router(config)#lpts pifib hardware police location 0/2/CPU0
RP/0/0/CPU0:router(config-pifib-policer-per-node)#flow bgp configured rate 4000
RP/0/0/CPU0:router(config-pifib-policer-per-node)#flow bgp known rate 4000
RP/0/0/CPU0:router(config-pifib-policer-per-node)#flow bgp default rate 4000
RP/0/0/CPU0:router(config-pifib-policer-per-node)#commit
```

## Logging Neighbor Changes

Logging neighbor changes is enabled by default. Use the **log neighbor changes disable** command to turn off logging. The **no log neighbor changes disable** command can also be used to turn logging back on if it has been disabled.

## Adjusting BGP Timers

Perform this task to set the timers for BGP neighbors.

BGP uses certain timers to control periodic activities, such as the sending of keepalive messages and the interval after which a neighbor is assumed to be down if no messages are received from the neighbor during the interval. The values set using the **timers bgp** command in router configuration mode can be overridden on particular neighbors using the **timers** command in the neighbor configuration mode.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **timers bgp** *keepalive hold-time*
4. **neighbor** *ip-address*
5. **timers** *keepalive hold-time*
6. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 123	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>timers bgp</b> <i>keepalive hold-time</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# timers bgp 30 90	Sets a default keepalive time and a default hold time for all neighbors.
<b>Step 4</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 5</b>	<b>timers</b> <i>keepalive hold-time</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# timers 60 220	(Optional) Sets the keepalive timer and the hold-time timer for the BGP neighbor.
<b>Step 6</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# end or RP/0/0/CPU0:router(config-bgp-nbr)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> </ul> </li> </ul>



	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Changing the BGP Default Local Preference Value

Perform this task to set the default local preference value for BGP paths.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp default local-preference** *value*
4. Do one of the following:
  - **end**
  - **commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# <code>configure</code>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# <code>router bgp 120</code>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>bgp default local-preference</b> <i>value</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# <code>bgp default local-preference 200</code>	Sets the default local preference value from the default of 100, making it either a more preferable path (over 100) or less preferable path (under 100).
<b>Step 4</b>	Do one of the following:	Saves configuration changes.

	Command or Action	Purpose
	<ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# end OR RP/0/0/CPU0:router(config-bgp)# commit</pre>	<ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <ul style="list-style-type: none"> <li>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</li> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring the MED Metric for BGP

Perform this task 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).

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **default-metric** *value*
4. Do one of the following:
  - **end**
  - **commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> <pre>RP/0/0/CPU0:router# configure</pre>	Enters global configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config)# router bgp 120</pre>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<p><b>default-metric</b> <i>value</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# default metric 10</pre>	Sets the default metric, which is used to set the MED to advertise to peers for routes that do not already have a metric set (routes that were received with no MED attribute).
<b>Step 4</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# end OR RP/0/0/CPU0:router(config-bgp)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Weights

Perform this task to assign a weight to routes received from a neighbor. A weight is a number that you can assign to a path so that you can control the best-path selection process. If you have particular neighbors that you want to prefer for most of your traffic, you can use the **weight** command to assign a higher weight to all routes learned from that neighbor.

### Before You Begin



#### Note

The **clear bgp** command must be used for the newly configured weight to take effect.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **address-family** { **ipv4** | **ipv6** } **unicast**
6. **weight** *weight-value*
7. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 4</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 5</b>	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either the IPv4 or IPv6 address family and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).

	Command or Action	Purpose
<b>Step 6</b>	<p><b>weight</b> <i>weight-value</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# weight 41150</pre>	<p>Assigns a weight to all routes learned through the neighbor.</p>
<b>Step 7</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/0/CPU0:router(config-bgp-nbr-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Tuning the BGP Best-Path Calculation

Perform this task to change the default BGP best-path calculation behavior.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp bestpath med missing-as-worst**
4. **bgp bestpath med always**
5. **bgp bestpath med confed**
6. **bgp bestpath as-path ignore**
7. **bgp bestpath compare-routerid**
8. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 126	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>bgp bestpath med missing-as-worst</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp bestpath med missing-as-worst	Directs the BGP software to consider a missing MED attribute in a path as having a value of infinity, making this path the least desirable path.
<b>Step 4</b>	<b>bgp bestpath med always</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp bestpath med always	Configures the BGP speaker in the specified autonomous system to compare MEDs among all the paths for the prefix, regardless of the autonomous system from which the paths are received.
<b>Step 5</b>	<b>bgp bestpath med confed</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp bestpath med confed	Enables BGP software to compare MED values for paths learned from confederation peers.

	Command or Action	Purpose
<b>Step 6</b>	<p><b>bgp bestpath as-path ignore</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# bgp bestpath as-path ignore</pre>	Configures the BGP software to ignore the autonomous system length when performing best-path selection.
<b>Step 7</b>	<p><b>bgp bestpath compare-routerid</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# bgp bestpath compare-routerid</pre>	Configure the BGP speaker in the autonomous system to compare the router IDs of similar paths.
<b>Step 8</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# end or RP/0/0/CPU0:router(config-bgp)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Indicating BGP Back-door Routes

Perform this task 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.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { *ipv4* | *ipv6* } **unicast**
4. **network** { *ip-address / prefix-length* | *ip-address mask* } **backdoor**
5. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either the IPv4 or IPv6 address family and enters address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 4</b>	<b>network</b> { <i>ip-address / prefix-length</i>   <i>ip-address mask</i> } <b>backdoor</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# network 172.20.0.0/16	Configures the local router to originate and advertise the specified network.
<b>Step 5</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# end	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:                 Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:             </li> </ul>



	Command or Action	Purpose
	or  <pre>RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> <ul style="list-style-type: none"> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Aggregate Addresses

Perform this task to create aggregate entries in a BGP routing table.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { *ipv4* | *ipv6* } **unicast**
4. **aggregate-address** *address/mask-length* [ *as-set* ] [ *as-confed-set* ] [ *summary-only* ] [ *route-policy route-policy-name* ]
5. Do one of the following:
  - **end**
  - **commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b>  <pre>RP/0/0/CPU0:router# configure</pre>	Enters global configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config)# router bgp 120</pre>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<p><b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	<p>Specifies either the IPv4 or IPv6 address family and enters address family configuration submenu.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
<b>Step 4</b>	<p><b>aggregate-address</b> <i>address/mask-length</i> [ <b>as-set</b> ] [ <b>as-confed-set</b> ] [ <b>summary-only</b> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# aggregate-address 10.0.0.0/8 as-set</pre>	<p>Creates 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.</p> <ul style="list-style-type: none"> <li>• The <b>as-set</b> keyword generates autonomous system set path information and community information from contributing paths.</li> <li>• The <b>as-confed-set</b> keyword generates autonomous system confederation set path information from contributing paths.</li> <li>• The <b>summary-only</b> keyword filters all more specific routes from updates.</li> <li>• The <b>route-policy</b> <i>route-policy-name</i> keyword and argument specify the route policy used to set the attributes of the aggregate route.</li> </ul>
<b>Step 5</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# end or RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistributing iBGP Routes into IGP

Perform this task to redistribute iBGP routes into an Interior Gateway Protocol (IGP), such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF).



**Note** 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.



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

### SUMMARY STEPS

1. **configure**
2. **router bgp *as-number***
3. **bgp redistribute-internal**
4. Do one of the following:
  - **end**
  - **commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>bgp redistribute-internal</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp redistribute-internal	Allows the redistribution of iBGP routes into an IGP, such as IS-IS or OSPF.
<b>Step 4</b>	Do one of the following:	Saves configuration changes.

Command or Action	Purpose
<ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp) # end or RP/0/0/CPU0:router(config-bgp) # commit</pre>	<ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <ul style="list-style-type: none"> <li>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</li> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistributing Prefixes into Multiprotocol BGP

Perform this task to redistribute prefixes from another protocol into multiprotocol BGP.

Redistribution is the process of injecting prefixes from one routing protocol into another routing protocol. This task shows how to inject prefixes from another routing protocol into multiprotocol BGP. Specifically, prefixes that are redistributed into multiprotocol BGP using the **redistribute** command are injected into the unicast database, the multicast database, or both.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **ipv4** | **ipv6** } **unicast**
4. Do one of the following:
  - **redistribute connected** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute eigrp** *process-id* [ **match** { **external** | **internal** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute isis** *process-id* [ **level** { **1** | **1-inter-area** | **2** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospf** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospfv3** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute rip** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute static** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
5. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either the IPv4 or IPv6 address family and enters address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).

	Command or Action	Purpose
<b>Step 4</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>redistribute connected</b> [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute eigrp</b> <i>process-id</i> [ <b>match</b> { <b>external</b>   <b>internal</b> } ] [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute isis</b> <i>process-id</i> [ <b>level</b> { <b>1</b>   <b>1-inter-area</b>   <b>2</b> } ] [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute ospf</b> <i>process-id</i> [ <b>match</b> { <b>external</b> [ <b>1</b>   <b>2</b> ]   <b>internal</b>   <b>nssa-external</b> [ <b>1</b>   <b>2</b> ] } ] [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute ospfv3</b> <i>process-id</i> [ <b>match</b> { <b>external</b> [ <b>1</b>   <b>2</b> ]   <b>internal</b>   <b>nssa-external</b> [ <b>1</b>   <b>2</b> ] } ] [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute rip</b> [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute static</b> [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# redistribute ospf 110</pre>	<p>Causes routes from the specified instance to be redistributed into BGP.</p>
<b>Step 5</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# end or RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <ul style="list-style-type: none"> <li>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</li> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"><li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li></ul>

## Configuring BGP Route Dampening

Perform this task to configure and monitor BGP route dampening.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **ipv4** | **ipv6** } **unicast**
4. **bgp dampening** [ *half-life* [ *reuse suppress max-suppress-time* ] ] | **route-policy** *route-policy-name* ]
5. Do one of the following:
  - **end**
  - **commit**
6. **show bgp** [ **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **tunnel** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** [ **rd** *rd-address* ] | **vrf** { *vrf-name* | **all** } [ **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** ] | **vpn6 unicast** [ **rd** *rd-address* ] ] **flap-statistics**
7. **show bgp** [ **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** [ **rd** *rd-address* ] | **vrf** { *vrf-name* | **all** } [ **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** ] | **vpn6 unicast** [ **rd** *rd-address* ] ] **flap-statistics** **regex** *regular-expression*
8. **show bgp** [ **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled** -**unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** [ **rd** *rd-address* ] | **vrf** { *vrf-name* | **all** } [ **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** ] | **vpn6 unicast** [ **rd** *rd-address* ] ] **flap-statistics** **route-policy** *route-policy-name*
9. **show bgp** [ **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled** -**unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** [ **rd** *rd-address* ] | **vrf** { *vrf-name* | **all** } [ **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** ] | **vpn6 unicast** [ **rd** *rd-address* ] ] **flap-statistics** { *ip-address* { *mask* | */prefix-length* } }
10. **show bgp** [ **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** [ **rd** *rd-address* ] | **vrf** { *vrf-name* | **all** } [ **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** ] | **vpn6 unicast** [ **rd** *rd-address* ] ] **flap-statistics** { *ip-address* [ { *mask* | */prefix-length* } ] [ **longer-prefixes** ] ] }
11. **clear bgp** { **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** | **vrf** { *vrf-name* | **all** } { **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** } | **vpn6 unicast** } **flap-statistics**
12. **clear bgp** { **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** | **vrf** { *vrf-name* | **all** } { **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** } | **vpn6 unicast** } **flap-statistics** **regex** *regular-expression*
13. **clear bgp** { **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** | **vrf** { *vrf-name* | **all** } { **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** } | **vpn6 unicast** } **flap-statistics** **route-policy** *route-policy-name*
14. **clear bgp** { **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **vpn4 unicast** |



```
vrf { vrf-name | all } { ipv4 { unicast | labeled-unicast } | ipv6 unicast } | vpnv6 unicast }
flap-statistics network / mask-length
```

15. **clear bgp** { ipv4 { unicast | multicast | labeled-unicast | all } | ipv6 { unicast | multicast | all | labeled-unicast } | all { unicast | multicast | all | labeled-unicast } | vpnv4 unicast | vrf { vrf-name | all } { ipv4 { unicast | labeled-unicast } | ipv6 unicast } | vpnv6 unicast } flap-statistics ip-address / mask-length

16. **show bgp** [ ipv4 { unicast | multicast | labeled-unicast | all } | ipv6 { unicast | multicast | all | labeled-unicast } | all { unicast | multicast | all | labeled-unicast } | vpnv4 unicast [ rd rd-address ] | vrf { vrf-name | all } [ ipv4 { unicast | labeled-unicast } | ipv6 unicast ] | vpnv6 unicast [ rd rd-address ] ] dampened-paths

17. **clear bgp** { ipv4 { unicast | multicast | labeled-unicast | all } | ipv6 { unicast | multicast | all | labeled-unicast } | all { unicast | multicast | all | labeled-unicast } | vpnv4 unicast | vrf { vrf-name | all } { ipv4 { unicast | labeled-unicast } | ipv6 unicast } | vpnv6 unicast } dampening [ ip-address / mask-length ]

**DETAILED STEPS**

	Command or Action	Purpose
Step 1	<p><b>configure</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# configure</pre>	Enters global configuration mode.
Step 2	<p><b>router bgp</b> as-number</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config)# router bgp 120</pre>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<p><b>address-family</b> { ipv4   ipv6 } unicast</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	<p>Specifies either the IPv4 or IPv6 address family and enters address family configuration submenu.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
Step 4	<p><b>bgp dampening</b> [ half-life [ reuse suppress max-suppress-time ]   route-policy route-policy-name ]</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# bgp dampening 30 1500 10000 120</pre>	Configures BGP dampening for the specified address family.
Step 5	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• end</li> <li>• commit</li> </ul>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:</li> </ul> <pre>Uncommitted changes found, commit</pre>

	Command or Action	Purpose
	<p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# end or RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<pre>them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> <ul style="list-style-type: none"> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
<b>Step 6</b>	<pre>show bgp [ ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   tunnel }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast [ rd rd-address ]   vrf { vrf-name   all } [ ipv4 { unicast   labeled-unicast }   ipv6 unicast ]   vpnv6 unicast [ rd rd-address ] ] flap-statistics</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show bgp flap statistics</pre>	Displays BGP flap statistics.
<b>Step 7</b>	<pre>show bgp [ ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   labeled-unicast }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast [ rd rd-address ]   vrf { vrf-name   all } [ ipv4 { unicast   labeled-unicast }   ipv6 unicast ]   vpnv6 unicast [ rd rd-address ] ] flap-statistics regexp regular-expression</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show bgp flap-statistics regexp _1\$</pre>	Displays BGP flap statistics for all paths that match the regular expression.
<b>Step 8</b>	<pre>show bgp [ ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   labeled -unicast }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast [ rd rd-address ]   vrf { vrf-name   all } [ ipv4 { unicast   labeled-unicast }   ipv6</pre>	Displays BGP flap statistics for the specified route policy.

	Command or Action	Purpose
	<p><b>unicast</b> ]   <b>vpn6 unicast</b> [ <b>rd rd-address</b> ] ] <b>flap-statistics</b>  <b>route-policy route-policy-name</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config)# show bgp flap-statistics route-policy policy_A</pre>	
<b>Step 9</b>	<p><b>show bgp</b> [ <b>ipv4 { unicast   multicast   labeled-unicast   all }</b>   <b>ipv6 { unicast   multicast   all   labeled-unicast }</b>   <b>all { unicast   multicast   all   labeled-unicast }</b> ]   <b>vpn4 unicast</b> [ <b>rd rd-address</b> ]   <b>vrf { vrf-name   all }</b> [ <b>ipv4 { unicast   labeled-unicast }</b>   <b>ipv6 unicast</b> ]   <b>vpn6 unicast</b> [ <b>rd rd-address</b> ] ] <b>flap-statistics</b> { <b>ip-address { mask   /prefix-length }</b> }</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show bgp flap-statistics 172.20.1.1</pre>	Displays BGP flap for the specified prefix.
<b>Step 10</b>	<p><b>show bgp</b> [ <b>ipv4 { unicast   multicast   labeled-unicast   all }</b>   <b>ipv6 { unicast   multicast   all   labeled-unicast }</b>   <b>all { unicast   multicast   all   labeled-unicast }</b> ]   <b>vpn4 unicast</b> [ <b>rd rd-address</b> ]   <b>vrf { vrf-name   all }</b> [ <b>ipv4 { unicast   labeled-unicast }</b> ]   <b>ipv6 unicast</b> ]   <b>vpn6 unicast</b> [ <b>rd rd-address</b> ] ] <b>flap-statistics</b> { <b>ip-address [ { mask   /prefix-length } [ longer-prefixes ] ]</b> }</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show bgp flap-statistics 172.20.1.1 longer-prefixes</pre>	Displays BGP flap statistics for more specific entries for the specified IP address.
<b>Step 11</b>	<p><b>clear bgp</b> { <b>ipv4 { unicast   multicast   labeled-unicast   all }</b>   <b>ipv6 { unicast   multicast   all   labeled-unicast }</b>   <b>all { unicast   multicast   all   labeled-unicast }</b> ]   <b>vpn4 unicast</b>   <b>vrf { vrf-name   all }</b> } { <b>ipv4 { unicast   labeled-unicast }</b>   <b>ipv6 unicast</b> }   <b>vpn6 unicast</b> } <b>flap-statistics</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp all all flap-statistics</pre>	Clears BGP flap statistics for all routes.
<b>Step 12</b>	<p><b>clear bgp</b> { <b>ipv4 { unicast   multicast   labeled-unicast   all }</b>   <b>ipv6 { unicast   multicast   all   labeled-unicast }</b>   <b>all { unicast   multicast   all   labeled-unicast }</b> ]   <b>vpn4 unicast</b>   <b>vrf { vrf-name   all }</b> } { <b>ipv4 { unicast   labeled-unicast }</b>   <b>ipv6 unicast</b> }   <b>vpn6 unicast</b> } <b>flap-statistics regexp regular-expression</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp ipv4 unicast flap-statistics regexp _1\$</pre>	Clears BGP flap statistics for all paths that match the specified regular expression.

	Command or Action	Purpose
Step 13	<pre>clear bgp { ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   labeled-unicast }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast   vrf { vrf-name   all } { ipv4 { unicast   labeled-unicast }   ipv6 unicast }   vpnv6 unicast } flap-statistics route-policy route-policy-name</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp ipv4 unicast flap-statistics route-policy policy_A</pre>	Clears BGP flap statistics for the specified route policy.
Step 14	<pre>clear bgp { ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   labeled-unicast }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast   vrf { vrf-name   all } { ipv4 { unicast   labeled-unicast }   ipv6 unicast }   vpnv6 unicast } flap-statistics network / mask-length</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp ipv4 unicast flap-statistics 192.168.40.0/24</pre>	Clears BGP flap statistics for the specified network.
Step 15	<pre>clear bgp { ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   labeled-unicast }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast   vrf { vrf-name   all } { ipv4 { unicast   labeled-unicast }   ipv6 unicast }   vpnv6 unicast } flap-statistics ip-address / mask-length</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp ipv4 unicast flap-statistics 172.20.1.1</pre>	Clears BGP flap statistics for routes received from the specified neighbor.
Step 16	<pre>show bgp [ ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   labeled -unicast }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast [ rd rd-address ]   vrf { vrf-name   all } [ ipv4 { unicast   labeled-unicast }   ipv6 unicast ]   vpnv6 unicast [ rd rd-address ]] dampened-paths</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show bgp dampened paths</pre>	Displays the dampened routes, including the time remaining before they are unsuppressed.
Step 17	<pre>clear bgp { ipv4 { unicast   multicast   labeled-unicast   all }   ipv6 { unicast   multicast   all   labeled-unicast }   all { unicast   multicast   all   labeled-unicast }   vpnv4 unicast   vrf { vrf-name</pre>	Clears route dampening information and unsuppresses the suppressed routes.

	Command or Action	Purpose
	<p>  all } { ipv4 { unicast   labeled-unicast }   ipv6 unicast }   vpv6 unicast } dampening [ ip-address / mask-length ]</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp dampening</pre>	<p><b>Caution</b> Always use the <b>clear bgp dampening</b> command for an individual address-family. The <b>all</b> option for address-families with <b>clear bgp dampening</b> should never be used during normal functioning of the system. For example, use</p> <pre>clear bgp ipv4 unicast dampening prefix x.x.x./y</pre>

## Applying Policy When Updating the Routing Table

Perform this task to apply a routing policy to routes being installed into the routing table.

### Before You Begin

See the *Implementing Routing Policy on Cisco IOS XR Software* module of *Cisco IOS XR Routing Configuration Guide for the Cisco XR 12000 Series Router* (this publication) for a list of the supported attributes and operations that are valid for table policy filtering.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { ipv4 | ipv6 } **unicast**
4. **table-policy** *policy-name*
5. Do one of the following:
  - **end**
  - **commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<p><b>configure</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# configure</pre>	Enters global configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config)# router bgp 120.6</pre>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<p><b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	Specifies either the IPv4 or IPv6 address family and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 4</b>	<p><b>table-policy</b> <i>policy-name</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# table-policy tbl-plcy-A</pre>	Applies the specified policy to routes being installed into the routing table.
<b>Step 5</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# end or RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Setting BGP Administrative Distance

Perform this task to specify the use of administrative distances that can be used to prefer one class of route over another.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **ipv4** | **ipv6** } **unicast**
4. **distance bgp** *external-distance internal-distance local-distance*
5. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 4</b>	<b>distance bgp</b> <i>external-distance internal-distance local-distance</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# distance bgp 20 20 200	Sets the external, internal, and local administrative distances to prefer one class of routes over another. The higher the value, the lower the trust rating.
<b>Step 5</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# end	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:               <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> </li> </ul>

Command or Action	Purpose
or  RP/0/0/CPU0:router(config-bgp-af)# commit	<ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> <ul style="list-style-type: none"> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring a BGP Neighbor Group and Neighbors

Perform this task to configure BGP neighbor groups and apply the neighbor group configuration to a neighbor. A neighbor group is a template that holds address family-independent and address family-dependent configurations associated with the neighbor.

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-dependent configurations. The inherited configuration can be overridden if you directly configure commands for the neighbor or configure session groups or address family groups through the **use** command.

You can configure an address family-independent configuration under the neighbor group. An address family-dependent configuration requires you to configure the address family under the neighbor group to enter address family submode.

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

After specifying the neighbor group name using the **neighbor group** command, you can assign options to the neighbor group.



### Note

All commands that can be configured under a specified neighbor group can be configured under a neighbor.



## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **ipv4** | **ipv6** } **unicast**
4. **exit**
5. **neighbor-group** *name*
6. **remote-as** *as-number*
7. **address-family** { **ipv4** | **ipv6** } **unicast**
8. **route-policy** *route-policy-name* { **in** | **out** }
9. **exit**
10. **exit**
11. **neighbor** *ip-address*
12. **use neighbor-group** *group-name*
13. **remote-as** *as-number*
14. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 4</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.

	Command or Action	Purpose
<b>Step 5</b>	<b>neighbor-group</b> <i>name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor-group nbr-grp-A	Places the router in neighbor group configuration mode.
<b>Step 6</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 7</b>	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 8</b>	<b>route-policy</b> <i>route-policy-name</i> { <b>in</b>   <b>out</b> }  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# route-policy drop-as-1234 in	(Optional) Applies the specified policy to inbound IPv4 unicast routes.
<b>Step 9</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp-af)# exit	Exits the current configuration mode.
<b>Step 10</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp)# exit	Exits the current configuration mode.
<b>Step 11</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 12</b>	<b>use neighbor-group</b> <i>group-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# use neighbor-group nbr-grp-A	(Optional) Specifies that the BGP neighbor inherit configuration from the specified neighbor group.

	Command or Action	Purpose
<b>Step 13</b>	<p><b>remote-as</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2002</pre>	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 14</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# end or RP/0/0/CPU0:router(config-bgp-nbr)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring a Route Reflector for BGP

Perform this task to configure a route reflector for BGP.

All the neighbors configured with the **route-reflector-client** command are members of the client group, and the remaining iBGP peers are members of the nonclient group for the local route reflector.

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 can 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. The **bgp cluster-id** command is used to configure the cluster ID when the cluster has more than one route reflector.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp cluster-id** *cluster-id*
4. **neighbor** *ip-address*
5. **remote-as** *as-number*
6. **address-family** { **ipv4** | **ipv6** } **unicast**
7. **route-reflector-client**
8. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>bgp cluster-id</b> <i>cluster-id</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp cluster-id 192.168.70.1	Configures the local router as one of the route reflectors serving the cluster. It is configured with a specified cluster ID to identify the cluster.
<b>Step 4</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 5</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2003	Creates a neighbor and assigns a remote autonomous system number to it.

	Command or Action	Purpose
<b>Step 6</b>	<p><b>address-family { ipv4   ipv6 } unicast</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	<p>Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
<b>Step 7</b>	<p><b>route-reflector-client</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# route-reflector-client</pre>	<p>Configures the router as a BGP route reflector and configures the neighbor as its client.</p>
<b>Step 8</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# end Or RP/0/0/CPU0:router(config-bgp-nbr-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Route Filtering by Route Policy

Perform this task to configure BGP routing filtering by route policy.

### Before You Begin

See the *Implementing Routing Policy on Cisco IOS XR Software* module of *Cisco Cisco IOS XR Routing Configuration Guide* (this publication) for a list of the supported attributes and operations that are valid for inbound and outbound neighbor policy filtering.

## SUMMARY STEPS

1. **configure**
2. **route-policy** *name*
3. **end-policy**
4. **router bgp** *as-number*
5. **neighbor** *ip-address*
6. **address-family** { **ipv4** | **ipv6** } **unicast**
7. **route-policy** *route-policy-name* { **in** | **out** }
8. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>route-policy</b> <i>name</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# route-policy drop-as-1234 RP/0/0/CPU0:router(config-rpl)# if as-path passes-through '1234' then RP/0/0/CPU0:router(config-rpl)# apply check-communities RP/0/0/CPU0:router(config-rpl)# else RP/0/0/CPU0:router(config-rpl)# pass RP/0/0/CPU0:router(config-rpl)# endif	(Optional) Creates a route policy and enters route policy configuration mode, where you can define the route policy.
<b>Step 3</b>	<b>end-policy</b>  <b>Example:</b> RP/0/0/CPU0:router(config-rpl)# end-policy	(Optional) Ends the definition of a route policy and exits route policy configuration mode.
<b>Step 4</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.

	Command or Action	Purpose
<b>Step 5</b>	<p><b>neighbor</b> <i>ip-address</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24</pre>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 6</b>	<p><b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	<p>Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
<b>Step 7</b>	<p><b>route-policy</b> <i>route-policy-name</i> { <b>in</b>   <b>out</b> }</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# route-policy drop-as-1234 in</pre>	Applies the specified policy to inbound routes.
<b>Step 8</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/0/CPU0:router(config-bgp-nbr-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Next-Hop Trigger Delay

Perform this task to configure BGP next-hop trigger delay. The Routing Information Base (RIB) classifies the dampening notifications based on the severity of the changes. Event notifications are classified as critical

and noncritical. This task allows you to specify the minimum batching interval for the critical and noncritical events.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **ipv4** | **ipv6** } **unicast**
4. **nexthop trigger-delay** { **critical** *delay* | **non-critical** *delay* }
5. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# <code>configure</code>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# <code>router bgp 120</code>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# <code>address-family ipv4 unicast</code>	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 4</b>	<b>nexthop trigger-delay</b> { <b>critical</b> <i>delay</i>   <b>non-critical</b> <i>delay</i> }  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# <code>nexthop trigger-delay critical 15000</code>	Sets the critical next-hop trigger delay.
<b>Step 5</b>	Do one of the following:  • <b>end</b>  • <b>commit</b>	Saves configuration changes.  • When you issue the <b>end</b> command, the system prompts you to commit changes:  <code>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</code>



	Command or Action	Purpose
	<p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# end or RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> <ul style="list-style-type: none"> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Disabling Next-Hop Processing on BGP Updates

Perform this task to disable next-hop calculation for a neighbor and insert your own address in the next-hop field of BGP updates. Disabling the calculation of the best next hop to use when advertising a route causes all routes to be advertised with the network device as the next hop.



### Note

Next-hop processing can be disabled for address family group, neighbor group, or neighbor address family.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **address-family** { **ipv4** | **ipv6** } **unicast**
6. **next-hop-self**
7. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 4</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 206	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 5</b>	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 6</b>	<b>next-hop-self</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# next-hop-self	Sets the next-hop attribute for all routes advertised to the specified neighbor to the address of the local router. Disabling the calculation of the best next hop to use when advertising a route causes all routes to be advertised with the local network device as the next hop.
<b>Step 7</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# end or RP/0/0/CPU0:router(config-bgp-nbr-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre> Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]: </pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP Community and Extended-Community Advertisements

Perform this task to specify that community attributes should be sent to an eBGP neighbor.

Perform this task to specify that community/extended-community attributes should be sent to an eBGP neighbor. These attributes are not sent to an eBGP neighbor by default. By contrast, they are always sent to iBGP neighbors. This section provides examples on how to enable sending community attributes. The **send-community-ebgp** keyword can be replaced by the **send-extended-community-ebgp** keyword to enable sending extended-communities.



### Note

If the **send-community-ebgp** 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.



### Note

BGP community and extended-community filtering cannot be configured for iBGP neighbors. Communities and extended-communities are always sent to iBGP neighbors

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **address-family** { *ipv4* | *ipv6* } **unicast**
6. **send-community-ebgp**
7. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 4</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 5</b>	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast	Enters neighbor address family configuration mode for the specified address family.
<b>Step 6</b>	<b>send-community-ebgp</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# send-community-ebgp	Specifies that the router send community attributes (which are disabled by default for eBGP neighbors) to a specified eBGP neighbor.
<b>Step 7</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/0/CPU0:router(config-bgp-nbr-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"><li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li><li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li><li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li></ul>

## Configuring the BGP Cost Community

Perform this task to configure the BGP cost community.

BGP receives multiple paths to the same destination and it uses the best-path algorithm to decide which is the best path to install in RIB. To enable users to determine an exit point after partial comparison, the cost community is defined to tie-break equal paths during the best-path selection process.

## SUMMARY STEPS

1. **configure**
2. **route-policy** *name*
3. **set extcommunity cost** { *cost-extcommunity-set-name* | *cost-inline-extcommunity-set* } [ **additive** ]
4. **end-policy**
5. **router bgp** *as-number*
6. Do one of the following:
  - **default-information originate**
  - **aggregate-address** *address/mask-length* [ *as-set* ] [ *as-confed-set* ] [ **summary-only** ] [ **route-policy** *route-policy-name* ]
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **redistribute connected** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **redistribute eigrp** *process-id* [ **match** { *external* | *internal* } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **redistribute isis** *process-id* [ **level** { *1* | *1-inter-area* | *2* } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **redistribute ospf** *process-id* [ **match** { *external* [ *1* | *2* ] | *internal* | *nssa-external* [ *1* | *2* ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
7. Do one of the following:
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **redistribute ospfv3** *process-id* [ **match** { *external* [ *1* | *2* ] | *internal* | *nssa-external* [ *1* | *2* ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **redistribute rip** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **redistribute static** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *vpn4 unicast* | *vpn6 unicast* } **network** { *ip-address/prefix-length* | *ip-address mask* } [ **route-policy** *route-policy-name* ]
  - **neighbor** *ip-address* **remote-as** *as-number* **address-family** { *ipv4 unicast* | *ipv4 multicast* | *ipv4 labeled-unicast* | *ipv4 tunnel* | *ipv4 mdt* | *ipv6 unicast* | *ipv6 multicast* | *ipv6 labeled-unicast* | *vpn4 unicast* | *vpn6 unicast* }

- **route-policy** *route-policy-name* { **in** | **out** }

8. Do one of the following:

- **end**
- **commit**

9. **show bgp** [ **vrf** *vrf-name* ] *ip-address*

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>route-policy</b> <i>name</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# route-policy costA	Enters route policy configuration mode and specifies the name of the route policy to be configured.
<b>Step 3</b>	<b>set extcommunity cost</b> { <i>cost-extcommunity-set-name</i>   <i>cost-inline-extcommunity-set</i> } [ <b>additive</b> ]  <b>Example:</b> RP/0/0/CPU0:router(config)# set extcommunity cost cost_A	Specifies the BGP extended community attribute for cost.
<b>Step 4</b>	<b>end-policy</b>  <b>Example:</b> RP/0/0/CPU0:router(config)# end-policy	Ends the definition of a route policy and exits route policy configuration mode.
<b>Step 5</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode allowing you to configure the BGP routing process.
<b>Step 6</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>default-information originate</b></li> <li>• <b>aggregate-address</b> <i>address/mask-length</i> [ <b>as-set</b> ] [ <b>as-confed-set</b> ] [ <b>summary-only</b> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>address-family</b> { <b>ipv4 unicast</b>   <b>ipv4 multicast</b>   <b>ipv4 tunnel</b>   <b>ipv4 mdt</b>   <b>ipv6 unicast</b>   <b>ipv6 multicast</b>   <b>vpn4 unicast</b>   <b>vpn6 unicast</b> } <b>redistribute connected</b> [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> </ul>	Applies the cost community to the attach point (route policy).

	Command or Action	Purpose
	<ul style="list-style-type: none"> <li>• <code>address-family { ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast } redistribute eigrp process-id [ match { external   internal } ] [ metric metric-value ] [ route-policy route-policy-name ]</code></li> <li>• <code>address-family { ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast } redistribute isis process-id [ level { 1   1-inter-area   2 } ] [ metric metric-value ] [ route-policy route-policy-name ]</code></li> <li>• <code>address-family { ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast } redistribute ospf process-id [ match { external [ 1   2 ]   internal   nssa-external [ 1   2 ] } ] [ metric metric-value ] [ route-policy route-policy-name ]</code></li> </ul>	
<b>Step 7</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <code>address-family { ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast } redistribute ospfv3 process-id [ match { external [ 1   2 ]   internal   nssa-external [ 1   2 ] } ] [ metric metric-value ] [ route-policy route-policy-name ]</code></li> <li>• <code>address-family { ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast } redistribute rip [ metric metric-value ] [ route-policy route-policy-name ]</code></li> <li>• <code>address-family { ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast } redistribute static [ metric metric-value ] [ route-policy route-policy-name ]</code></li> <li>• <code>address-family { ipv4 unicast   ipv4 multicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   vpv4 unicast   vpv6 unicast } network { ip-address/prefix-length   ip-address mask } [ route-policy route-policy-name ]</code></li> <li>• <code>neighbor ip-address remote-as as-number address-family { ipv4 unicast   ipv4 multicast   ipv4 labeled-unicast   ipv4 tunnel   ipv4 mdt   ipv6 unicast   ipv6 multicast   ipv6 labeled-unicast   vpv4 unicast   vpv6 unicast }</code></li> <li>• <code>route-policy route-policy-name { in   out }</code></li> </ul>	
<b>Step 8</b>	Do one of the following:	Saves configuration changes.



	Command or Action	Purpose
	<ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# end OR RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>
<b>Step 9</b>	<pre>show bgp [ vrf vrf-name ] ip-address</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show bgp 172.168.40.24</pre>	<p>Displays the cost community in the following format:</p> <p>Cost: <i>POI</i> : <i>cost-community-ID</i> : <i>cost-number</i></p>

## Configuring Software to Store Updates from a Neighbor

Perform this task to configure the software to store updates received from a neighbor.

The **soft-reconfiguration inbound** command causes a route refresh request to be sent to the neighbor if the neighbor is route refresh capable. If the neighbor is not route refresh capable, the neighbor must be reset to relearn received routes using the **clear bgp soft** command. See the [Resetting Neighbors Using BGP Inbound Soft Reset](#), page 138.

**Note**

Storing updates from a neighbor works only if either the neighbor is route refresh capable or the **soft-reconfiguration inbound** command is configured. Even if the neighbor is route refresh capable and the **soft-reconfiguration inbound** command is configured, the original routes are not stored unless the **always** option is used with the command. The original routes can be easily retrieved with a route refresh request. Route refresh sends a request to the peer to resend its routing information. The **soft-reconfiguration inbound** command stores all paths received from the peer in an unmodified form and refers to these stored paths during the clear. Soft reconfiguration is memory intensive.

**SUMMARY STEPS**

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **address-family** { **ipv4** | **ipv6** } **unicast**
5. **soft-reconfiguration inbound** [ **always**]
6. Do one of the following:
  - **end**
  - **commit**

**DETAILED STEPS**

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# <code>configure</code>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# <code>router bgp 120</code>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# <code>neighbor 172.168.40.24</code>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.

	Command or Action	Purpose
<b>Step 4</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	<p>Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
<b>Step 5</b>	<b>soft-reconfiguration inbound [ always]</b>  <b>Example:</b> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound always</pre>	<p>Configures the software to store updates received from a specified neighbor. Soft reconfiguration inbound causes the software to store the original unmodified route in addition to a route that is modified or filtered. This allows a “soft clear” to be performed after the inbound policy is changed.</p> <p>Soft reconfiguration enables the software to store the incoming updates before apply policy if route refresh is not supported by the peer (otherwise a copy of the update is not stored). The <b>always</b> keyword forces the software to store a copy even when route refresh is supported by the peer.</p>
<b>Step 6</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/0/CPU0:router(config-bgp-nbr-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Distributed BGP

Perform this task to configure distributed BGP. Configuring distributed BGP includes starting the speaker process and allocating the speaker process to a neighbor.

**Before You Begin**

**Note** If BGP is running in standalone mode, the **clear bgp current-mode** or **clear bgp vrf all \*** command must be used to switch from standalone mode to distributed mode.

**SUMMARY STEPS**

1. **configure**
2. **router bgp** *as-number*
3. **distributed speaker** *id*
4. **commit**
5. **address-family** { **ipv4** | **ipv6** } **unicast**
6. **exit**
7. **neighbor** *ip-address*
8. **remote-as** *as-number*
9. **speaker-id** *id*
10. **address-family** { **ipv4** | **ipv6** } **unicast**
11. **end**
12. **clear bgp current-mode**

**DETAILED STEPS**

	<b>Command or Action</b>	<b>Purpose</b>
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>distributed speaker</b> <i>id</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# distributed speaker 2	Specifies the speaker process to start.
<b>Step 4</b>	<b>commit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# commit	Saves the configuration changes to the running configuration file and remains within the configuration session.

	Command or Action	Purpose
<b>Step 5</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 6</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits address family mode.
<b>Step 7</b>	<b>neighbor ip-address</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 8</b>	<b>remote-as as-number</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2002	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 9</b>	<b>speaker-id id</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# speaker-id 2	Allocates a neighbor to a specified speaker process.
<b>Step 10</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 11</b>	<b>end</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# end	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre> Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]: </pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul>
<b>Step 12</b>	<b>clear bgp current-mode</b>  <b>Example:</b> RP/0/0/CPU0:router# clear bgp current-mode	Switches from standalone mode to distributed mode.

## Configuring a VPN Routing and Forwarding Instance in BGP

Layer 2 and Layer 3 (virtual private network) VPN can be configured only if there is an available Layer 3 VPN license for the line card slot on which the feature is being configured.

If the advanced IP license is enabled, 4096 Layer 3 VPN routing and forwarding instances (VRFs) can be configured on an interface. If the infrastructure VRF license is enabled, eight Layer 3 VRFs can be configured on the line card. See the Software Entitlement on Cisco IOS XR Software module in *Cisco IOS XR System Management Configuration Guide for the Cisco XR 12000 Series Router* for more information on advanced IP licencing.

The following error message appears if the appropriate licence is not enabled:

```
RP/0/0/CPU0:router#LC/0/0/CPU0:Dec 15 17:57:53.653 : rsi_agent[247]:
%LICENSE-ASR9K_LICENSE-2-INFRA_VRF_NEEDED : 5 VRF(s) are configured without license
A9K-iVRF-LIC in violation of the Software Right To Use Agreement.
This feature may be disabled by the system without the appropriate license.
Contact Cisco to purchase the license immediately to avoid potential service interruption.
```

The following tasks are used to configure a VPN routing and forwarding (VRF) instance in BGP:

### Defining the Virtual Routing and Forwarding Tables in Provider Edge Routers

Perform this task to define the VPN routing and forwarding (VRF) tables in the provider edge (PE) routers.

## SUMMARY STEPS

1. **configure**
2. **vrf** *vrf-name*
3. **address-family** { **ipv4** | **ipv6** } **unicast**
4. **maximum prefix** *maximum* [ *threshold* ]
5. **import route-policy** *policy-name*
6. **import route-target** [ *as-number : nn* | *ip-address : nn* ]
7. **export route-policy** *policy-name*
8. **export route-target** [ *as-number : nn* | *ip-address : nn* ]
9. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# vrf vrf_pe	Configures a VRF instance.
<b>Step 3</b>	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either the IPv4 or IPv6 address family and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 4</b>	<b>maximum prefix</b> <i>maximum</i> [ <i>threshold</i> ]  <b>Example:</b> RP/0/0/CPU0:router(config-vrf-af)# maximum prefix 2300	Configures a limit to the number of prefixes allowed in a VRF table.  A maximum number of routes is applicable only to dynamic routing protocols and not to static or connected routes.  You can specify a threshold percentage of the prefix limit using the <i>mid-threshold</i> argument.
<b>Step 5</b>	<b>import route-policy</b> <i>policy-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf-af)# import route-policy policy_a	(Optional) Provides finer control over what gets imported into a VRF. This import filter discards prefixes that do not match the specified <i>policy-name</i> argument.

	Command or Action	Purpose
<b>Step 6</b>	<b>import route-target</b> [ <i>as-number : nn</i>   <i>ip-address : nn</i> ]  <b>Example:</b> RP/0/0/CPU0:router(config-vrf-af)# import route-target 234:222	Specifies a list of route target (RT) extended communities. Only prefixes that are associated with the specified import route target extended communities are imported into the VRF.
<b>Step 7</b>	<b>export route-policy</b> <i>policy-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf-af)# export route-policy policy_b	(Optional) Provides finer control over what gets exported into a VRF. This export filter discards prefixes that do not match the specified <i>policy-name</i> argument.
<b>Step 8</b>	<b>export route-target</b> [ <i>as-number : nn</i>   <i>ip-address : nn</i> ]  <b>Example:</b> RP/0/0/CPU0:router(config-vrf-af)# export route-target 123;234	Specifies a list of route target extended communities. Export route target communities are associated with prefixes when they are advertised to remote PEs. The remote PEs import them into VRFs which have import RTs that match these exported route target communities.
<b>Step 9</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-vrf-af)# end or RP/0/0/CPU0:router(config-vrf-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring the Route Distinguisher

The route distinguisher (RD) makes prefixes unique across multiple VPN routing and forwarding (VRF) instances.



In the L3VPN multipath same route distinguisher (RD) environment, the determination of whether to install a prefix in RIB or not is based on the prefix's bestpath. In a rare misconfiguration situation, where the best path is not a valid path to be installed in RIB, BGP drops the prefix and does not consider the other paths. The behavior is different for different RD setup, where the non-best multipath will be installed if the best multipath is invalid to be installed in RIB.

Perform this task to configure the RD.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp router-id** *ip-address*
4. **vrf** *vrf-name*
5. **rd** { *as-number : nn* | *ip-address : nn* | **auto** }
6. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Enters BGP configuration mode allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>bgp router-id</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp router-id 10.0.0.0	Configures a fixed router ID for the BGP-speaking router.
<b>Step 4</b>	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# vrf vrf_pe	Configures a VRF instance.
<b>Step 5</b>	<b>rd</b> { <i>as-number : nn</i>   <i>ip-address : nn</i>   <b>auto</b> } }	Configures the route distinguisher.

	Command or Action	Purpose
	<p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf)# rd 345:567</pre>	<p>Use the <b>auto</b> keyword if you want the router to automatically assign a unique RD to the VRF.</p> <p>Automatic assignment of RDs is possible only if a router ID is configured using the <b>bgp router-id</b> command in router configuration mode. This allows you to configure a globally unique router ID that can be used for automatic RD generation. The router ID for the VRF does not need to be globally unique, and using the VRF router ID would be incorrect for automatic RD generation. Having a single router ID also helps in checkpointing RD information for BGP graceful restart, because it is expected to be stable across reboots.</p>
<p><b>Step 6</b></p>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf)# end OR RP/0/0/CPU0:router(config-bgp-vrf)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP to Advertise VRF Routes for Multicast VPN from PE to PE

Perform these tasks to enable multicast VPN routing for IPv4 and IPv6 address families from one provider edge (PE) router to another:

## Advertising VRF Routes for MVPNv4 from PE to PE

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp router-id** *ip-address*
4. **address-family** { **ipv4** | **ipv6** } **unicast**
5. **exit**
6. **address-family** **vpn4 unicast**
7. **exit**
8. **address-family** **ipv4 mdt**
9. **exit**
10. **neighbor** *ip-address*
11. **remote-as** *as-number*
12. **update-source** *type interface-path-id*
13. **address-family** { **ipv4** | **ipv6** } **unicast**
14. **exit**
15. **address-family** **vpn4 unicast**
16. **exit**
17. **vrf** *vrf-name*
18. **rd** { *as-number : nn* | *ip-address : nn* | **auto** }
19. **address-family** { **ipv4** | **ipv6** } **unicast**
20. Do one of the following:
  - **redistribute connected** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute eigrp** *process-id* [ **match** { **external** | **internal** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute isis** *process-id* [ **level** { **1** | **1-inter-area** | **2** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospf** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospfv3** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute rip** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute static** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
21. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 100	Enters BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>bgp router-id <i>ip-address</i></b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp router-id 1.1.1.1	Configures a fixed router ID for a BGP-speaking router.
<b>Step 4</b>	<b>address-family { <i>ipv4</i>   <i>ipv6</i> } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 5</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits IPv4 address family configuration submode and reenters BGP configuration submode.
<b>Step 6</b>	<b>address-family <i>vpn</i>v4 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family vpnv4 unicast	Enters VPNv4 address family configuration submode.
<b>Step 7</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits IPv4 address-family configuration submode and reenters BGP configuration submode.

	Command or Action	Purpose
<b>Step 8</b>	<b>address-family ipv4 mdt</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family ipv4 mdt	Configures an IPv4 address-family multicast distribution tree (MDT).
<b>Step 9</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
<b>Step 10</b>	<b>neighbor ip-address</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.16.1.1	Places the PE router in neighbor configuration submode.
<b>Step 11</b>	<b>remote-as as-number</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 100	Creates a neighbor and assigns the neighbor a remote autonomous system number, which can be from 1 to 65535.
<b>Step 12</b>	<b>update-source type interface-path-id</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# update-source loopback 0	<p>Allows sessions to use the primary IP address from a specific interface as the local address when forming a session with a neighbor.</p> <p>The <i>interface-type interface-id</i> arguments specify the type and ID number of the interface, such as GigabitEthernet or Loopback. Use the CLI help (?) to see a list of all the possible interface types and their ID numbers.</p>
<b>Step 13</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	<p>Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
<b>Step 14</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# exit	Exits the neighbor address family configuration submode.
<b>Step 15</b>	<b>address-family vpnv4 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# address-family vpnv4 unicast	<p>Specifies the address family as VPNv4 and enters address family configuration submode.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>

	Command or Action	Purpose
<b>Step 16</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# exit	Exits BGP neighbor address family configuration submode.
<b>Step 17</b>	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# vrf vpn1	Enables BGP routing for a particular VRF on the PE router.
<b>Step 18</b>	<b>rd { as-number : nn   ip-address : nn   auto }</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# rd 1:1	Configures the route distinguisher. <ul style="list-style-type: none"> <li>• Use the <b>auto</b> keyword if you want the router to automatically assign a unique RD to the VRF.</li> <li>• Automatic assignment of RDs is possible only if a router ID is configured using the <b>bgp router-id</b> command in router configuration mode. This allows you to configure a globally unique router ID that can be used for automatic RD generation.</li> </ul> <p>The router ID for the VRF does not need to be globally unique, and using the VRF router ID would be incorrect for automatic RD generation. Having a single router ID also helps in checkpointing RD information for BGP graceful restart, because it is expected to be stable across reboots.</p>
<b>Step 19</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 20</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>redistribute connected [ metric metric-value ] [ route-policy route-policy-name ]</b></li> <li>• <b>redistribute eigrp process-id [ match { external   internal } ] [ metric metric-value ] [ route-policy route-policy-name ]</b></li> <li>• <b>redistribute isis process-id [ level { 1   1-inter-area   2 } ] [ metric metric-value ] [ route-policy route-policy-name ]</b></li> <li>• <b>redistribute ospf process-id [ match { external [ 1   2 ]   internal   nssa-external</b></li> </ul>	Configures redistribution of a protocol into the VRF address family context.

	Command or Action	Purpose
	<p>[ 1   2 ]} [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</p> <ul style="list-style-type: none"> <li>• redistribute ospfv3 <i>process-id</i> [ match { external [ 1   2 ]   internal   nssa-external [ 1   2 ] } ] [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</li> <li>• redistribute rip [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</li> <li>• redistribute static [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf-af)# redistribute ospf 1</pre>	
<p><b>Step 21</b></p>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• end</li> <li>• commit</li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf-af)# end OR RP/0/0/CPU0:router(config-bgp-vrf-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:             <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Advertising VRF Routes for MVPNv6 from PE to PE

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **bgp router-id** *ip-address*
4. **address-family ipv6 unicast**
5. **address-family vpnv6 unicast**
6. **exit**
7. **neighbor-group** *vrf-name*
8. **remote-as** *as-number*
9. **update-source** *interface-type interface-id*
10. **address-family vpnv6 unicast**
11. **exit**
12. **exit**
13. **neighbor** *ip-address*
14. **remote-as** *as-number*
15. **use neighbor-group** *vpn-name*
16. **update-source** *interface-type interface-id*
17. **address-family ipv6 unicast**
18. **exit**
19. **address-family vpnv6 unicast**
20. **exit**
21. **exit**
22. **vrf** *vrf-name*
23. **rd** { *as-number : nn* | *ip-address : nn* | **auto** }
24. **exit**
25. **vrf** *vrf-name*
26. **rd** { *as-number : nn* | *ip-address : nn* | **auto** }
27. **address-family ipv6 unicast**
28. Do one of the following:
  - **end**
  - **commit**



## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 100	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>bgp router-id</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# bgp router-id 1.1.1.1 Configures a fixed router ID for a BGP-speaking router.	
Step 4	<b>address-family ipv6 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family ipv6 unicast	Specifies the address family as IPv6 and enters IPv6 neighbor address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
Step 5	<b>address-family vpnv6 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family vpnv6 unicast	Enters VPNv6 address family configuration submode.
Step 6	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the VPNv6 address family configuration submode.
Step 7	<b>neighbor-group</b> <i>vrf-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor-group vpn22	Places the router in neighbor group configuration submode.
Step 8	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp)# remote-as 100	Creates a neighbor and assigns the neighbor a remote autonomous system number, which can be from 1 to 65535.

	Command or Action	Purpose
<b>Step 9</b>	<b>update-source</b> <i>interface-type interface-id</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr) # update-source loopback 0	Allows sessions to use the primary IP address from a specific interface as the local address when forming a session with a neighbor.  The <i>interface-type interface-id</i> arguments specify the type and ID number of the interface, such as ATM, POS, Loopback. Use the CLI help (?) to see a list of all the possible interface types and their ID numbers.
<b>Step 10</b>	<b>address-family</b> <b>vpn6 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp) # address-family vpn6 unicast	Specifies the address family as VPNv6 and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 11</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp-af) # exit	Exits the neighbor group address family configuration submode.
<b>Step 12</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbrgrp) # exit	Exits BGP neighbor group configuration submode.
<b>Step 13</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp) # neighbor 1.1.1.2	Places a PE router in neighbor group configuration submode.
<b>Step 14</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr) # remote-as 100	Creates a neighbor and assigns it a remote autonomous system number, which can be from 1 to 65535.
<b>Step 15</b>	<b>use neighbor-group</b> <i>vpn-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr) # use neighbor-group vpn22	(Optional) Specifies that the BGP neighbor inherits the configuration from the specified VPN neighbor group.
<b>Step 16</b>	<b>update-source</b> <i>interface-type interface-id</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr) # update-source loopback 0	Allows sessions to use the primary IP address from a specific interface as the local address when forming a session with a neighbor.  The <i>interface-type interface-id</i> arguments specify the type and ID number of the interface, such as ATM, POS, Loopback. Use the

	Command or Action	Purpose
		CLI help (?) to see a list of all the possible interface types and their ID numbers.
<b>Step 17</b>	<b>address-family ipv6 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# address-family ipv6 unicast	Specifies the address family as IPv6 and enters IPv6 neighbor address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 18</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# exit	Exits BGP neighbor address family configuration submenu.
<b>Step 19</b>	<b>address-family vpnv6 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# address-family vpnv6 unicast	Specifies the address family as VPNv6 and enters address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 20</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# exit	Exits the neighbor address family configuration submenu.
<b>Step 21</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# exit	Exits the BGP neighbor configuration submenu.
<b>Step 22</b>	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# vrf vpn1	Enters BGP VRF configuration submenu.
<b>Step 23</b>	<b>rd { as-number : nn   ip-address : nn   auto }</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# rd 111:1	Configures the route distinguisher. <ul style="list-style-type: none"> <li>• Use the <b>auto</b> keyword if you want the router to automatically assign a unique RD to the VRF.</li> <li>• Automatic assignment of RDs is possible only if a router ID is configured using the <b>bgp router-id</b> command in router configuration mode. This allows you to configure a globally unique router ID that can be used for automatic RD generation.</li> </ul> <p>The router ID for the VRF does not need to be globally unique, and using the VRF router ID would be incorrect for automatic RD generation. Having a single router ID also helps in checkpointing RD information for BGP graceful restart, because it is expected to be stable across reboots.</p>

	Command or Action	Purpose
<b>Step 24</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# exit	Exits BGP VRF configuration submode.
<b>Step 25</b>	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# vrf vpn1	Enables BGP routing for a particular VRF on the PE router.
<b>Step 26</b>	<b>rd { as-number : nn   ip-address : nn   auto }</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# rd 1:1	Configures the route distinguisher. <ul style="list-style-type: none"> <li>• Use the <b>auto</b> keyword if you want the router to automatically assign a unique RD to the VRF.</li> <li>• Automatic assignment of RDs is possible only if a router ID is configured using the <b>bgp router-id</b> command in router configuration mode. This allows you to configure a globally unique router ID that can be used for automatic RD generation. The router ID for the VRF does not need to be globally unique, and using the VRF router ID would be incorrect for automatic RD generation. Having a single router ID also helps in checkpointing RD information for BGP graceful restart, because it is expected to be stable across reboots.</li> </ul>
<b>Step 27</b>	<b>address-family ipv6 unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# address-family ipv6 unicast	Specifies the address family as IPv6 and enters IPv6 VRF address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 28</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf-af)# end or RP/0/0/CPU0:router(config-bgp-vrf-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:   <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring PE-PE or PE-RR Interior BGP Sessions

To enable BGP to carry VPN reachability information between provider edge (PE) routers you must configure the PE-PE interior BGP (iBGP) sessions. A PE uses VPN information carried from the remote PE router to determine VPN connectivity and the label value to be used so the remote (egress) router can demultiplex the packet to the correct VPN during packet forwarding.

The PE-PE, PE-route reflector (RR) iBGP sessions are defined to all PE and RR routers that participate in the VPNs configured in the PE router.

Perform this task to configure PE-PE iBGP sessions and to configure global VPN options on a PE.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **vpn4 unicast** | **vpn6 unicast** }
4. **exit**
5. **neighbor** *ip-address*
6. **remote-as** *as-number*
7. **description** *text*
8. **password** { **clear** | **encrypted** } *password*
9. **shutdown**
10. **timers** *keepalive hold-time*
11. **update-source** *type interface-id*
12. **address-family** { **vpn4 unicast** | **vpn6 unicast** }
13. **route-policy** *route-policy-name* **in**
14. **route-policy** *route-policy-name* **out**
15. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family vpn4 unicast	Enters VPN address family configuration mode.
<b>Step 4</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
<b>Step 5</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.16.1.1	Configures a PE iBGP neighbor.
<b>Step 6</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 1	Assigns the neighbor a remote autonomous system number.
<b>Step 7</b>	<b>description</b> <i>text</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# description neighbor 172.16.1.1	(Optional) Provides a description of the neighbor. The description is used to save comments and does not affect software function.
<b>Step 8</b>	<b>password</b> { <b>clear</b>   <b>encrypted</b> } <i>password</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# password encrypted 123abc	Enables Message Digest 5 (MD5) authentication on the TCP connection between the two BGP neighbors.

	Command or Action	Purpose
Step 9	<p><b>shutdown</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# shutdown</pre>	Terminates any active sessions for the specified neighbor and removes all associated routing information.
Step 10	<p><b>timers</b> <i>keepalive hold-time</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# timers 12000 200</pre>	Set the timers for the BGP neighbor.
Step 11	<p><b>update-source</b> <i>type interface-id</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# update-source gigabitEthernet 0/1/5/0</pre>	Allows iBGP sessions to use the primary IP address from a specific interface as the local address when forming an iBGP session with a neighbor.
Step 12	<p><b>address-family</b> { <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }</p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# address-family vpn4 unicast</pre>	Enters VPN neighbor address family configuration mode.
Step 13	<p><b>route-policy</b> <i>route-policy-name in</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# route-policy pe-pe-vpn-in in</pre>	Specifies a routing policy for an inbound route. The policy can be used to filter routes or modify route attributes.
Step 14	<p><b>route-policy</b> <i>route-policy-name out</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# route-policy pe-pe-vpn-out out</pre>	Specifies a routing policy for an outbound route. The policy can be used to filter routes or modify route attributes.
Step 15	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/0/CPU0:router(config-bgp-nbr-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Route Reflector to Hold Routes That Have a Defined Set of RT Communities

A provider edge (PE) needs to hold the routes that match the import route targets (RTs) of the VPNs configured on it. The PE router can discard all other VPNv4 (Cisco XR 12000 Series Router and Cisco CRS-1) and VPNv6 (Cisco XR 12000 Series Router only) routes. But, a route reflector (RR) must retain all VPNv4 and VPNv6 routes, because it might peer with PE routers and different PEs might require different RT-tagged VPNv4 and VPNv6 routes (making RRs non-scalable). You can configure an RR to only hold routes that have a defined set of RT communities. Also, a number of the RRs can be configured to service a different set of VPNs (thereby achieving some scalability). A PE is then made to peer with all RRs that service the VRFs 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 refreshes to the RRs and retrieves the relevant VPN routes.



### Note

Note that this process can be more efficient if the PE-RR session supports extended community outbound route filter (ORF).

Perform this task to configure a reflector to retain routes tagged with specific RTs.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **vpn4 unicast** | **vpn6 unicast** }
4. **retain route-target** { **all** | **route-policy** *route-policy-name* }
5. Do one of the following:
  - **end**
  - **commit**



## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family</b> { <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family vpn4 unicast	Enters VPN address family configuration mode.
<b>Step 4</b>	<b>retain route-target</b> { <b>all</b>   <b>route-policy</b> <i>route-policy-name</i> }  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# retain route-target route-policy rr_ext-comm	Configures a reflector to retain routes tagged with particular RTs. Use the <i>route-policy-name</i> argument for the policy name that lists the extended communities that a path should have in order for the RR to retain that path.  <b>Note</b> The <b>all</b> keyword is not required, because this is the default behavior of a route reflector.
<b>Step 5</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr-af)# end OR RP/0/0/CPU0:router(config-bgp-nbr-af)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring BGP as a PE-CE Protocol

Perform this task to configure BGP on the PE and establish PE-CE communication using BGP.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **vrf** *vrf-name*
4. **bgp router-id** *ip-address*
5. **label-allocation-mode** { **per-ce** | **per-vrf** }
6. **address-family** { **ipv4** | **ipv6** } **unicast**
7. **network** { *ip-address / prefix-length* | *ip-address mask* }
8. **aggregate-address** *address / mask-length*
9. **exit**
10. **neighbor** *ip-address*
11. **remote-as** *as-number*
12. **password** { **clear** | **encrypted** } *password*
13. **ebgp-multihop** [*tth-value*]
14. Do one of the following:
  - **address-family** { **ipv4** | **ipv6** } **unicast**
  - **address-family** { **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** }
15. **site-of-origin** [*as-number : nn* | *ip-address : nn*]
16. **as-override**
17. **allowas-in** [*as-occurrence-number*]
18. **route-policy** *route-policy-name* **in**
19. **route-policy** *route-policy-name* **out**
20. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# vrf vrf_pe_2	Enables BGP routing for a particular VRF on the PE router.
Step 4	<b>bgp router-id</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# bgp router-id 172.16.9.9	Configures a fixed router ID for a BGP-speaking router.
Step 5	<b>label-allocation-mode</b> { <b>per-ce</b>   <b>per-vrf</b> }  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-ce	Configures the MPLS/VPN label allocation mode. <ul style="list-style-type: none"> <li>• The <b>per-ce</b> keyword configures the per-CE label allocation mode to avoid an extra lookup on the PE router and conserve label space (per-prefix is the default label allocation mode). In this mode, the PE router allocates one label for every immediate next-hop (in most cases, this would be a CE router). This 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 would be one for each CE rather than one for each VRF. Because BGP knows all the next hops, it assigns a label for each next hop (not for each PE-CE interface). When the outgoing interface is a multiaccess interface and the media access control (MAC) address of the neighbor is not known, Address Resolution Protocol (ARP) is triggered during packet forwarding.</li> <li>• The <b>per-vrf</b> keyword configures the same label to be used for all the routes advertised from a unique VRF.</li> </ul>

	Command or Action	Purpose
<b>Step 6</b>	<b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 7</b>	<b>network</b> { <i>ip-address / prefix-length</i>   <i>ip-address mask</i> }  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf-af)# network 172.16.5.5	Originates a network prefix in the address family table in the VRF context.
<b>Step 8</b>	<b>aggregate-address</b> <i>address / mask-length</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf-af)# aggregate-address 10.0.0.0/24	Configures aggregation in the VRF address family context to summarize routing information to reduce the state maintained in the core. This summarization introduces some inefficiency in the PE edge, because an additional lookup is required to determine the ultimate next hop for a packet. When configured, a summary prefix is advertised instead of a set of component prefixes, which are more specifics of the aggregate. The PE advertises only one label for the aggregate. Because component prefixes could have different next hops to CEs, an additional lookup has to be performed during data forwarding.
<b>Step 9</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf-af)# exit	Exits the current configuration mode.
<b>Step 10</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# neighbor 10.0.0.0	Configures a CE neighbor. The <i>ip-address</i> argument must be a private address.
<b>Step 11</b>	<b>remote-as</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf-nbr)# remote-as 2	Configures the remote AS for the CE neighbor.
<b>Step 12</b>	<b>password</b> { <i>clear</i>   <i>encrypted</i> } <i>password</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf-nbr)# password encrypted 234xyz	Enable Message Digest 5 (MD5) authentication on a TCP connection between two BGP neighbors.

	Command or Action	Purpose
<b>Step 13</b>	<b>ebgp-multihop</b> [ <i>ttl-value</i> ]  <b>Example:</b>  RP/0/0/CPU0:router(config-bgp-vrf-nbr)# ebgp-multihop 55	Configures the CE neighbor to accept and attempt BGP connections to external peers residing on networks that are not directly connected.
<b>Step 14</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>address-family { ipv4   ipv6 } unicast</b></li> <li>• <b>address-family {ipv4 {unicast   labeled-unicast}   ipv6 unicast}</b></li> </ul> <b>Example:</b>  RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 (unicast or labeled-unicast) or IPv6 unicast address family and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 15</b>	<b>site-of-origin</b> [ <i>as-number : nn</i>   <i>ip-address : nn</i> ]  <b>Example:</b>  RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# site-of-origin 234:111	Configures the site-of-origin (SoO) extended community. Routes that are learned from this CE neighbor are tagged with the SoO extended community before being advertised to the rest of the PEs. SoO is frequently used to detect loops when as-override is configured on the PE router. If the prefix is looped back to the same site, the PE detects this and does not send the update to the CE.
<b>Step 16</b>	<b>as-override</b>  <b>Example:</b>  RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# as-override	Configures AS override on the PE router. This causes the PE router to replace the CE's ASN with its own (PE) ASN.  <b>Note</b> This loss of information could lead to routing loops; to avoid loops caused by as-override, use it in conjunction with site-of-origin.
<b>Step 17</b>	<b>allows-in</b> [ <i>as-occurrence-number</i> ]  <b>Example:</b>  RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# allows-in 5	Allows an AS path with the PE autonomous system number (ASN) a specified number of times.  Hub and spoke VPN networks need the looping back of routing information to the HUB PE through the HUB CE. When this happens, due to the presence of the PE ASN, the looped-back information is dropped by the HUB PE. To avoid this, use the <b>allows-in</b> command to allow prefixes even if they have the PEs ASN up to the specified number of times.
<b>Step 18</b>	<b>route-policy</b> <i>route-policy-name</i> <b>in</b>  <b>Example:</b>  RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# route-policy pe_ce_in_policy in	Specifies a routing policy for an inbound route. The policy can be used to filter routes or modify route attributes.

	Command or Action	Purpose
<b>Step 19</b>	<p><b>route-policy</b> <i>route-policy-name</i> <b>out</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# route-policy pe_ce_out_policy out</pre>	Specifies a routing policy for an outbound route. The policy can be used to filter routes or modify route attributes.
<b>Step 20</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# end OR RP/0/0/CPU0:router(config-bgp-vrf-nbr-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Redistribution of IGP to BGP

Perform this task to configure redistribution of a protocol into the VRF address family.

Even if Interior Gateway Protocols (IGPs) are used as the PE-CE protocol, the import logic happens through BGP. Therefore, all IGP routes have to be imported into the BGP VRF table.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **vrf** *vrf-name*
4. **address-family** { **ipv4** | **ipv6** } **unicast**
5. Do one of the following:
  - **redistribute connected** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute eigrp** *process-id* [ **match** { **external** | **internal** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute isis** *process-id* [ **level** { **1** | **1-inter-area** | **2** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospf** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospfv3** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute rip** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute static** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
6. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>vrf</b> <i>vrf-name</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# vrf vrf_a	Enables BGP routing for a particular VRF on the PE router.

	Command or Action	Purpose
<b>Step 4</b>	<p><b>address-family { ipv4   ipv6 } unicast</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	<p>Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
<b>Step 5</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>redistribute connected [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</b></li> <li>• <b>redistribute eigrp <i>process-id</i> [ match { external   internal } ] [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</b></li> <li>• <b>redistribute isis <i>process-id</i> [ level { 1   1-inter-area   2 } ] [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</b></li> <li>• <b>redistribute ospf <i>process-id</i> [ match { external [ 1   2 ]   internal   nssa-external [ 1   2 ] } ] [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</b></li> <li>• <b>redistribute ospfv3 <i>process-id</i> [ match { external [ 1   2 ]   internal   nssa-external [ 1   2 ] } ] [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</b></li> <li>• <b>redistribute rip [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</b></li> <li>• <b>redistribute static [ metric <i>metric-value</i> ] [ route-policy <i>route-policy-name</i> ]</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf-af)# redistribute eigrp 23</pre>	<p>Configures redistribution of a protocol into the VRF address family context.</p> <p>The <b>redistribute</b> command is used if BGP is not used between the PE-CE routers. If BGP is used between PE-CE routers, the IGP that is used has to be redistributed into BGP to establish VPN connectivity with other PE sites. Redistribution is also required for inter-table import and export.</p>
<b>Step 6</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf-af)# end or RP/0/0/CPU0:router(config-bgp-vrf-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <ul style="list-style-type: none"> <li>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</li> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> </ul> </li> </ul>



	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Keychains for BGP

Keychains provide secure authentication by supporting different MAC authentication algorithms and provide graceful key rollover. Perform this task to configure keychains for BGP. This task is optional.



### Note

If a keychain is configured for a neighbor group or a session group, a neighbor using the group inherits the keychain. Values of commands configured specifically for a neighbor override inherited values.

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **remote-as** *as-number*
5. **keychain** *name*
6. Do one of the following:
  - **end**
  - **commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b>  RP/0/0/CPU0:router# <code>configure</code>	Enters global configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<p><b>router bgp</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config)# router bgp 120</pre>	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<p><b>neighbor</b> <i>ip-address</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24</pre>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 4</b>	<p><b>remote-as</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2002</pre>	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 5</b>	<p><b>keychain</b> <i>name</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# keychain kych_a</pre>	Configures keychain-based authentication.
<b>Step 6</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# end or RP/0/0/CPU0:router(config-bgp-nbr)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring an MDT Address Family Session in BGP

Perform this task to configure an IPv4 multicast distribution tree (MDT) subaddress family identifier (SAFI) session in BGP, which can also be used for MVPNv6 network distribution.

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **ipv4** | **ipv6** } **unicast**
4. **exit**
5. **address-family** { **vpn4** | **vpn6** } **unicast**
6. **exit**
7. **address-family ipv4 mdt**
8. **exit**
9. **neighbor** *ip-address*
10. **remote-as** *as-number*
11. **update-source** *interface-type interface-id*
12. **address-family** { **ipv4** | **ipv6** } **unicast**
13. **exit**
14. **address-family** {**vpn4** | **vpn6**} **unicast**
15. **exit**
16. **address-family ipv4 mdt**
17. **exit**
18. **vrf** *vrf-name*
19. **rd** { *as-number:nn* | *ip-address:nn* | **auto** }
20. **address-family** { **ipv4** | **ipv6** } **unicast**
21. Do one of the following:
  - **redistribute connected** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute eigrp** *process-id* [ **match** { **external** | **internal** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute isis** *process-id* [ **level** { **1** | **1-inter-area** | **2** } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospf** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute ospfv3** *process-id* [ **match** { **external** [ **1** | **2** ] | **internal** | **nssa-external** [ **1** | **2** ] } ] [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute rip** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
  - **redistribute static** [ **metric** *metric-value* ] [ **route-policy** *route-policy-name* ]
22. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
Step 2	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
Step 3	<b>address-family</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
Step 4	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
Step 5	<b>address-family</b> { <b>vpn4</b>   <b>vpn6</b> } <b>unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family vpn4 unicast	Specifies the address family and enters the address family configuration submode.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).  <b>Note</b> Required if you are configuring multicast MVPN. If configuring MVPNv6, use the <b>vpn6</b> keyword
Step 6	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
Step 7	<b>address-family ipv4 mdt</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family ipv4 mdt	Specifies the multicast distribution tree (MDT) address family.
Step 8	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.

	Command or Action	Purpose
<b>Step 9</b>	<p><b>neighbor</b> <i>ip-address</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24</pre>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 10</b>	<p><b>remote-as</b> <i>as-number</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# remote-as 2002</pre>	Creates a neighbor and assigns a remote autonomous system number to it.
<b>Step 11</b>	<p><b>update-source</b> <i>interface-type interface-id</i></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# update-source loopback 0</pre>	<p>Allows sessions to use the primary IP address from a specific interface as the local address when forming a session with a neighbor.</p> <p>The <i>interface-type interface-id</i> arguments specify the type and ID number of the interface, such as ATM, POS, Loopback. Use the CLI help (?) to see a list of all the possible interface types and their ID numbers.</p>
<b>Step 12</b>	<p><b>address-family</b> { <i>ipv4</i>   <i>ipv6</i> } <b>unicast</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast</pre>	<p>Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submode.</p> <p>To see a list of all the possible keywords and arguments for this command, use the CLI help (?).</p>
<b>Step 13</b>	<p><b>exit</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# exit</pre>	(Optional) Exits the current configuration mode.
<b>Step 14</b>	<p><b>address-family</b> { <i>vpn4</i>   <i>vpn6</i> } <b>unicast</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr)# address-family vpn4 unicast</pre>	<p>(Optional) Enters address family configuration submode for the specified address family.</p> <p><b>Note</b> Required if you are configuring multicast MVPN. If configuring MVPNv6, use the <b>vpn6</b> keyword.</p>
<b>Step 15</b>	<p><b>exit</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-nbr-af)# exit</pre>	Exits the current configuration mode.
<b>Step 16</b>	<p><b>address-family</b> <i>ipv4</i> <b>mdt</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# address-family ipv4 mdt</pre>	Specifies the multicast distribution tree (MDT) address family.

	Command or Action	Purpose
<b>Step 17</b>	<b>exit</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# exit	Exits the current configuration mode.
<b>Step 18</b>	<b>vrf vrf-name</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# vrf vpn1	(Optional) Enables BGP routing for a particular VRF on the PE router.  <b>Note</b> Required if you are configuring multicast MVPN.
<b>Step 19</b>	<b>rd { as-number:nn   ip-address:nn   auto }</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-vrf)# rd 1:1	(Optional) Configures the route distinguisher.  <ul style="list-style-type: none"> <li>• Use the <b>auto</b> keyword if you want the router to automatically assign a unique RD to the VRF.</li> <li>• Automatic assignment of RDs is possible only if a router ID is configured using the <b>bgp router-id</b> command in router configuration mode. This allows you to configure a globally unique router ID that can be used for automatic RD generation.</li> </ul> <p>The router ID for the VRF does not need to be globally unique, and using the VRF router ID would be incorrect for automatic RD generation. Having a single router ID also helps in checkpointing RD information for BGP graceful restart, because it is expected to be stable across reboots.</p> <b>Note</b> Required if you are configuring multicast MVPN.
<b>Step 20</b>	<b>address-family { ipv4   ipv6 } unicast</b>  <b>Example:</b> RP/0/0/CPU0:router(config-vrf)# address-family ipv4 unicast	Specifies either an IPv4 or IPv6 address family unicast and enters address family configuration submenu.  To see a list of all the possible keywords and arguments for this command, use the CLI help (?).
<b>Step 21</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>redistribute connected</b> [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute eigrp</b> <i>process-id</i> [ <b>match</b> { <b>external</b>   <b>internal</b> } ] [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute isis</b> <i>process-id</i> [ <b>level</b> { <b>1</b>   <b>1-inter-area</b>   <b>2</b> } ] [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute ospf</b> <i>process-id</i> [ <b>match</b> { <b>external</b> [ <b>1</b>   <b>2</b> ]   <b>internal</b>   <b>nssa-external</b></li> </ul>	(Optional) Configures redistribution of a protocol into the VRF address family context.  <b>Note</b> Required if you are configuring multicast MVPN.

	Command or Action	Purpose
	<pre>[ 1   2 ]} [ metric metric-value ] [ route-policy route-policy-name ]</pre> <ul style="list-style-type: none"> <li>• <b>redistribute ospfv3</b> <i>process-id</i> [ <b>match</b> { <b>external</b> [ 1   2 ]   <b>internal</b>   <b>nssa-external</b> [ 1   2 ] } ] [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute rip</b> [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> <li>• <b>redistribute static</b> [ <b>metric</b> <i>metric-value</i> ] [ <b>route-policy</b> <i>route-policy-name</i> ]</li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-vrf-af)# redistribute eigrp 23</pre>	
<b>Step 22</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# end OR RP/0/0/CPU0:router(config-bgp-af)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Disabling a BGP Neighbor

Perform this task to administratively shut down a neighbor session without removing the configuration.



## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **neighbor** *ip-address*
4. **shutdown**
5. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 127	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>neighbor</b> <i>ip-address</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# neighbor 172.168.40.24	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address as a BGP peer.
<b>Step 4</b>	<b>shutdown</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# shutdown	Disables all active sessions for the specified neighbor.
<b>Step 5</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp-nbr)# end	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:               <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel) ?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> </ul> </li> </ul>

	Command or Action	Purpose
	or  RP/0/0/CPU0:router(config-bgp-nbr)# commit	<ul style="list-style-type: none"> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> <ul style="list-style-type: none"> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Resetting Neighbors Using BGP Inbound Soft Reset

Perform this task to trigger an inbound soft reset of the specified address families for the specified group or neighbors. The group is specified by the *\**, *ip-address*, *as-number*, or **external** keywords and arguments.

Resetting neighbors is useful if you change the inbound policy for the neighbors or any other configuration that affects the sending or receiving of routing updates. If an inbound soft reset is triggered, BGP 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.

### SUMMARY STEPS

1. **show bgp neighbors**
2. **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 } { \* | ip-address | as as-number | external } soft [ in [ prefix-filter ] | out ]**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>show bgp neighbors</b>  <b>Example:</b>  RP/0/0/CPU0:router# show bgp neighbors	Verifies that received route refresh capability from the neighbor is enabled.
<b>Step 2</b>	<b>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</b>	Soft resets a BGP neighbor. <ul style="list-style-type: none"> <li>• The <i>*</i> keyword resets all BGP neighbors.</li> <li>• The <i>ip-address</i> argument specifies the address of the neighbor to be reset.</li> </ul>

Command or Action	Purpose
<b>unicast</b> }   <b>vpn6 unicast</b> } { *   <i>ip-address</i>   <b>as</b> <i>as-number</i>   <b>external</b> } <b>soft</b> [ <b>in</b> [ <i>prefix-filter</i> ]   <b>out</b> ]  <b>Example:</b>  RP/0/0/CPU0:router# clear bgp ipv4 unicast 10.0.0.1 soft in	<ul style="list-style-type: none"> <li>The <i>as-number</i> argument specifies that all neighbors that match the autonomous system number be reset.</li> <li>The <b>external</b> keyword specifies that all external neighbors are reset.</li> </ul>

## Resetting Neighbors Using BGP Outbound Soft Reset

Perform this task to trigger an outbound soft reset of the specified address families for the specified group or neighbors. The group is specified by the *\**, *ip-address*, *as-number*, or **external** keywords and arguments.

Resetting neighbors is useful if you change the 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.

To determine whether the neighbor has advertised the ROUTE\_REFRESH capability, use the **show bgp neighbors** command.

### SUMMARY STEPS

1. **show bgp neighbors**
2. **clear bgp** { **ipv4** { **unicast** | **multicast** | **labeled-unicast** | **all** | **tunnel** | } | **ipv6** { **unicast** | **multicast** | **all** | **labeled-unicast** } | **all** { **unicast** | **multicast** | **all** | **labeled-unicast** | **mdt** | **tunnel** } | **vpn4 unicast** | **vrf** { *vrf-name* | **all** } { **ipv4** { **unicast** | **labeled-unicast** } | **ipv6 unicast** } | **vpn6 unicast** } { \* | *ip-address* | **as** *as-number* | **external** } **soft out**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>show bgp neighbors</b>  <b>Example:</b>  RP/0/0/CPU0:router# show bgp neighbors	Verifies that received route refresh capability from the neighbor is enabled.
<b>Step 2</b>	<b>clear bgp</b> { <b>ipv4</b> { <b>unicast</b>   <b>multicast</b>   <b>labeled-unicast</b>   <b>all</b>   <b>tunnel</b>   }   <b>ipv6</b> { <b>unicast</b>   <b>multicast</b>   <b>all</b>   <b>labeled-unicast</b> }   <b>all</b> { <b>unicast</b>   <b>multicast</b>   <b>all</b>   <b>labeled-unicast</b>   <b>mdt</b>   <b>tunnel</b> }   <b>vpn4 unicast</b>   <b>vrf</b> { <i>vrf-name</i>   <b>all</b> } { <b>ipv4</b> { <b>unicast</b>   <b>labeled-unicast</b> }   <b>ipv6 unicast</b> }   <b>vpn6 unicast</b> } { *   <i>ip-address</i>   <b>as</b> <i>as-number</i>   <b>external</b> } <b>soft out</b>	Soft resets a BGP neighbor. <ul style="list-style-type: none"> <li>The <i>*</i> keyword resets all BGP neighbors.</li> <li>The <i>ip-address</i> argument specifies the address of the neighbor to be reset.</li> <li>The <i>as-number</i> argument specifies that all neighbors that match the autonomous system number be reset.</li> </ul>

	Command or Action	Purpose
	<p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp ipv4 unicast 10.0.0.2 soft out</pre>	<ul style="list-style-type: none"> <li>The <b>external</b> keyword specifies that all external neighbors are reset.</li> </ul>

## Resetting Neighbors Using BGP Hard Reset

Perform this task to reset neighbors using a hard reset. A hard reset 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.

### SUMMARY STEPS

- 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 } { \* | ip-address | as as-number | external } [ graceful ] soft [ in [ prefix-filter ] | out ]**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<pre>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 } { *   ip-address   as as-number   external } [ graceful ] soft [ in [ prefix-filter ]   out ]</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp ipv4 unicast 10.0.0.3 graceful soft out</pre>	<p>Clears a BGP neighbor.</p> <ul style="list-style-type: none"> <li>The <b>*</b> keyword resets all BGP neighbors.</li> <li>The <i>ip-address</i> argument specifies the address of the neighbor to be reset.</li> <li>The <i>as-number</i> argument specifies that all neighbors that match the autonomous system number be reset.</li> <li>The <b>external</b> keyword specifies that all external neighbors are reset.</li> </ul> <p>The <b>graceful</b> keyword specifies a graceful restart.</p>

## Clearing Caches, Tables, and Databases

Perform this task to remove all contents of a particular cache, table, or database. The **clear bgp** command resets 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. Clearing a cache, table, or database can become necessary when the contents of the particular structure have become, or are suspected to be, invalid.

## SUMMARY STEPS

1. `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 } ip-address`
2. `clear bgp external`
3. `clear bgp *`

## DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>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 } ip-address</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp ipv4 172.20.1.1</pre>	Clears a specified neighbor.
Step 2	<pre>clear bgp external</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp external</pre>	Clears all external peers.
Step 3	<pre>clear bgp *</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# clear bgp *</pre>	Clears all BGP neighbors.

## Displaying System and Network Statistics

Perform this task to display specific statistics, such as the contents of BGP routing tables, caches, and databases. Information provided can be used to determine resource usage and solve network problems. You can also display information about node reachability and discover the routing path that the packets of your device are taking through the network.

## SUMMARY STEPS

1. **show bgp cidr-only**
2. **show bgp community** *community-list* [ **exact-match** ]
3. **show bgp regexp** *regular-expression*
4. **show bgp**
5. **show bgp neighbors** *ip-address* [ **advertised-routes** | **dampened-routes** | **flap-statistics** | **performance-statistics** | **received** *prefix-filter* | **routes** ]
6. **show bgp paths**
7. **show bgp neighbor-group** *group-name* **configuration**
8. **show bgp summary**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>show bgp cidr-only</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp cidr-only	Displays routes with nonnatural network masks (classless interdomain routing [CIDR]) routes.
<b>Step 2</b>	<b>show bgp community</b> <i>community-list</i> [ <b>exact-match</b> ]  <b>Example:</b> RP/0/0/CPU0:router# show bgp community 1081:5 exact-match	Displays routes that match the specified BGP community.
<b>Step 3</b>	<b>show bgp regexp</b> <i>regular-expression</i>  <b>Example:</b> RP/0/0/CPU0:router# show bgp regexp "^3 "	Displays routes that match the specified autonomous system path regular expression.
<b>Step 4</b>	<b>show bgp</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp	Displays entries in the BGP routing table.
<b>Step 5</b>	<b>show bgp neighbors</b> <i>ip-address</i> [ <b>advertised-routes</b>   <b>dampened-routes</b>   <b>flap-statistics</b>   <b>performance-statistics</b>   <b>received</b> <i>prefix-filter</i>   <b>routes</b> ]  <b>Example:</b> RP/0/0/CPU0:router# show bgp neighbors 10.0.101.1	Displays information about the BGP connection to the specified neighbor. <ul style="list-style-type: none"> <li>• The <b>advertised-routes</b> keyword displays all routes the router advertised to the neighbor.</li> <li>• The <b>dampened-routes</b> keyword displays the dampened routes that are learned from the neighbor.</li> <li>• The <b>flap-statistics</b> keyword displays flap statistics of the routes learned from the neighbor.</li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>The <b>performance-statistics</b> keyword displays performance statistics relating to work done by the BGP process for this neighbor.</li> <li>The <b>received <i>prefix-filter</i></b> keyword and argument display the received prefix list filter.</li> <li>The <b>routes</b> keyword displays routes learned from the neighbor.</li> </ul>
<b>Step 6</b>	<b>show bgp paths</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp paths	Displays all BGP paths in the database.
<b>Step 7</b>	<b>show bgp neighbor-group <i>group-name</i> configuration</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp neighbor-group group_1 configuration	Displays the effective configuration for a specified neighbor group, including any configuration inherited by this neighbor group.
<b>Step 8</b>	<b>show bgp summary</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp summary	Displays the status of all BGP connections.

## Displaying BGP Process Information

Perform this task to display specific BGP process information.

### SUMMARY STEPS

1. **show bgp process**
2. **show bgp ipv4 unicast summary**
3. **show bgp vpnv4 unicast summary**
4. **show bgp vrf ( *vrf-name* | all }**
5. **show bgp process detail**
6. **show bgp summary**
7. **show placement program bgp**
8. **show placement program brib**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>show bgp process</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp process	Displays status and summary information for the 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.
<b>Step 2</b>	<b>show bgp ipv4 unicast summary</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp ipv4 unicast summary	Displays a summary of the neighbors for the IPv4 unicast address family.
<b>Step 3</b>	<b>show bgp vpnv4 unicast summary</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp vpnv4 unicast summary	Displays a summary of the neighbors for the VPNv4 unicast address family.
<b>Step 4</b>	<b>show bgp vrf ( vrf-name   all )</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp vrf vrf_A	Displays BGP VPN virtual routing and forwarding (VRF) information.
<b>Step 5</b>	<b>show bgp process detail</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp processes detail	Displays detailed process information including the memory used by each of various internal structure types.
<b>Step 6</b>	<b>show bgp summary</b>  <b>Example:</b> RP/0/0/CPU0:router# show bgp summary	Displays the status of all BGP connections.
<b>Step 7</b>	<b>show placement program bgp</b>  <b>Example:</b> RP/0/0/CPU0:router# show placement program bgp	Displays BGP program information. <ul style="list-style-type: none"> <li>• If a program is shown as having ‘rejected locations’ (for example, locations where program cannot be placed), the locations in question can be viewed using the <b>show placement program bgp</b> command.</li> <li>• If a program has been placed but not started, the amount of elapsed time since the program was placed is displayed in the Waiting to start column.</li> </ul>
<b>Step 8</b>	<b>show placement program brib</b>	Displays bRIB program information.



	Command or Action	Purpose
	<p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show placement program brib</pre>	<ul style="list-style-type: none"> <li>• If a program is shown as having ‘rejected locations’ (for example, locations where program cannot be placed), the locations in question can be viewed using the <b>show placement program bgp</b> command.</li> <li>• If a program has been placed but not started, the amount of elapsed time since the program was placed is displayed in the Waiting to start column.</li> </ul>

## Monitoring BGP Update Groups

This task displays information related to the processing of BGP update groups.

### SUMMARY STEPS

1. **show bgp [ ipv4 { unicast | multicast | labeled-unicast | all | tunnel | } | ipv6 { unicast | all | labeled-unicast } | all { unicast | multicast | all | mdt | labeled-unicast | 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 ]]**

### DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>show bgp [ ipv4 { unicast   multicast   labeled-unicast   all   tunnel   }   ipv6 { unicast   all   labeled-unicast }   all { unicast   multicast   all   mdt   labeled-unicast   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 ]]</pre> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router# show bgp update-group 0.0</pre>	<p>Displays information about BGP update groups.</p> <ul style="list-style-type: none"> <li>• The <i>ip-address</i> argument displays the update groups to which that neighbor belongs.</li> <li>• The <i>process-id.index</i> argument selects a particular update group to display and is specified as follows: process ID (dot) index. Process ID range is from 0 to 254. Index range is from 0 to 4294967295.</li> <li>• The <b>summary</b> keyword displays summary information for neighbors in a particular update group.</li> <li>• If no argument is specified, this command displays information for all update groups (for the specified address family).</li> <li>• The <b>performance-statistics</b> keyword displays performance statistics for an update group.</li> </ul>

## Configuring BGP Nonstop Routing

Perform this task to configure BGP Nonstop Routing (NSR).

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **nsr**
4. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 120	Specifies the BGP AS number, and enters the BGP configuration mode, for configuring BGP routing processes.
<b>Step 3</b>	<b>nsr</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# nsr	Activates BGP Nonstop routing.
<b>Step 4</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# end or RP/0/0/CPU0:router(config-bgp)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes:               <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Configuring Best-External Path Advertisement

Perform the following tasks to advertise the best-external path to the iBGP and route-reflector peers:

### SUMMARY STEPS

- configure**
- router bgp** *as-number*
- Do one of the following
  - address-family** { **vpn4 unicast** | **vpn6 unicast** }
  - vrfvrf-name**{**ipv4 unicast**|**ipv6 unicast**}
- advertise best-external**
- Do one of the following:
  - end**
  - commit**

### DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# <b>configure</b>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# <b>router bgp</b> 100	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	Do one of the following <ul style="list-style-type: none"> <li><b>address-family</b> { <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }</li> <li><b>vrfvrf-name</b>{<b>ipv4 unicast</b> <b>ipv6 unicast</b>}</li> </ul>	Specifies the address family or VRF address family and enters the address family or VRF address family configuration submode.

	Command or Action	Purpose
	<p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# address-family vpv4 unicast</pre>	
<b>Step 4</b>	<p><b>advertise best-external</b></p> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp-af)# advertise best-external</pre>	Advertise the best-external path to the iBGP and route-reflector peers.
<b>Step 5</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# end or RP/0/0/CPU0:router(config-bgp)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> <li>• Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session.</li> </ul>

## Installing Primary Backup Path

Perform the following tasks to install a backup path into the forwarding table and provide prefix independent convergence (PIC) in case of a PE-CE link failure:

## SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. Do one of the following
  - **address-family** { **vpn4 unicast** | **vpn6 unicast** }
  - **vrfvrf-name**{**ipv4 unicast**|**ipv6 unicast**}
4. **additional-paths install backup**
5. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 100	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	Do one of the following <ul style="list-style-type: none"> <li>• <b>address-family</b> { <b>vpn4 unicast</b>   <b>vpn6 unicast</b> }</li> <li>• <b>vrfvrf-name</b>{<b>ipv4 unicast</b> <b>ipv6 unicast</b>}</li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family vpn4 unicast	Specifies the address family or VRF address family and enters the address family or VRF address family configuration submode.
<b>Step 4</b>	<b>additional-paths install backup</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# additional-paths install backup	Installs a backup path into the forwarding table and provides prefix independent convergence (PIC) in case of a PE-CE link failure.

	Command or Action	Purpose
<b>Step 5</b>	<p>Do one of the following:</p> <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <p><b>Example:</b></p> <pre>RP/0/0/CPU0:router(config-bgp)# end or RP/0/0/CPU0:router(config-bgp)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <ul style="list-style-type: none"> <li>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</li> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> </ul> <p>Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session</p>

## Retaining Allocated Local Label for Primary Path

Perform the following tasks to retain the previously allocated local label for the primary path on the primary PE for some configurable time after reconvergence:

### SUMMARY STEPS

1. **configure**
2. **router bgp** *as-number*
3. **address-family** { **vpn4 unicast** | **vpn6 unicast** }
4. **retain local-label** *minutes*
5. Do one of the following:
  - **end**
  - **commit**

## DETAILED STEPS

	Command or Action	Purpose
<b>Step 1</b>	<b>configure</b>  <b>Example:</b> RP/0/0/CPU0:router# configure	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> RP/0/0/CPU0:router(config)# router bgp 100	Specifies the autonomous system number and enters the BGP configuration mode, allowing you to configure the BGP routing process.
<b>Step 3</b>	<b>address-family { vpv4 unicast   vpv6 unicast }</b>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# address-family vpv4 unicast	Specifies the address family and enters the address family configuration submode.
<b>Step 4</b>	<b>retain local-label</b> <i>minutes</i>  <b>Example:</b> RP/0/0/CPU0:router(config-bgp-af)# retain local-label 10	Retains the previously allocated local label for the primary path on the primary PE for 10 minutes after reconvergence.
<b>Step 5</b>	Do one of the following: <ul style="list-style-type: none"> <li>• <b>end</b></li> <li>• <b>commit</b></li> </ul> <b>Example:</b> RP/0/0/CPU0:router(config-bgp)# end or RP/0/0/CPU0:router(config-bgp)# commit	Saves configuration changes. <ul style="list-style-type: none"> <li>• When you issue the <b>end</b> command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)?[cancel]:</pre> <ul style="list-style-type: none"> <li>◦ Entering <b>yes</b> saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.</li> <li>◦ Entering <b>no</b> exits the configuration session and returns the router to EXEC mode without committing the configuration changes.</li> <li>◦ Entering <b>cancel</b> leaves the router in the current configuration session without exiting or committing the configuration changes.</li> </ul> </li> </ul> Use the <b>commit</b> command to save the configuration changes to the running configuration file and remain within the configuration session

# Configuration Examples for Implementing BGP

This section provides the following configuration examples:

## Enabling BGP: Example

The following shows how to enable BGP.

```

prefix-set static
  2020::/64,
  2012::/64,
  10.10.0.0/16,
  10.2.0.0/24
end-set

route-policy pass-all
  pass
end-policy
route-policy set_next_hop_agg_v4
  set next-hop 10.0.0.1
end-policy

route-policy set_next_hop_static_v4
  if (destination in static) then
    set next-hop 10.1.0.1
  else
    drop
  endif
end-policy

route-policy set_next_hop_agg_v6
  set next-hop 2003::121
end-policy

route-policy set_next_hop_static_v6
  if (destination in static) then
    set next-hop 2011::121
  else
    drop
  endif
end-policy

router bgp 65000
  bgp fast-external-fallover disable
  bgp confederation peers
    65001
    65002
  bgp confederation identifier 1
  bgp router-id 1.1.1.1
  address-family ipv4 unicast
    aggregate-address 10.2.0.0/24 route-policy set_next_hop_agg_v4
    aggregate-address 10.3.0.0/24
    redistribute static route-policy set_next_hop_static_v4
  address-family ipv4 multicast
    aggregate-address 10.2.0.0/24 route-policy set_next_hop_agg_v4
    aggregate-address 10.3.0.0/24
    redistribute static route-policy set_next_hop_static_v4
  address-family ipv6 unicast
    aggregate-address 2012::/64 route-policy set_next_hop_agg_v6
    aggregate-address 2013::/64
    redistribute static route-policy set_next_hop_static_v6
  address-family ipv6 multicast
    aggregate-address 2012::/64 route-policy set_next_hop_agg_v6
    aggregate-address 2013::/64
    redistribute static route-policy set_next_hop_static_v6
  neighbor 10.0.101.60

```



```

remote-as 65000
address-family ipv4 unicast
address-family ipv4 multicast
neighbor 10.0.101.61
remote-as 65000
address-family ipv4 unicast
address-family ipv4 multicast
neighbor 10.0.101.62
remote-as 3
address-family ipv4 unicast
route-policy pass-all in
route-policy pass-all out
address-family ipv4 multicast
route-policy pass-all in
route-policy pass-all out
neighbor 10.0.101.64
remote-as 5
update-source Loopback0
address-family ipv4 unicast
route-policy pass-all in
route-policy pass-all out
address-family ipv4 multicast
route-policy pass-all in
route-policy pass-all out

```

## Displaying BGP Update Groups: Example

The following is sample output from the `show bgp update-group` command run in EXEC mode:

```

RP/0/0/CPU0:router# show bgp update-group

Update group for IPv4 Unicast, index 0.1:
Attributes:
  Outbound Route map:rm
  Minimum advertisement interval:30
  Messages formatted:2, replicated:2
  Neighbors in this update group:
    10.0.101.92

Update group for IPv4 Unicast, index 0.2:
Attributes:
  Minimum advertisement interval:30
  Messages formatted:2, replicated:2
  Neighbors in this update group:
    10.0.101.91

```

## BGP Neighbor Configuration: Example

The following example shows how BGP neighbors on an autonomous system are configured to share information. In the example, a BGP router is assigned to autonomous system 109, and two networks are listed as originating in the autonomous system. Then the addresses of three remote routers (and their autonomous systems) are listed. The router being configured shares information about networks 131. 108.0.0 and 192. 31.7.0 with the neighbor routers. The first router listed is in a different autonomous system; the second **neighbor** and **remote-as** commands specify an internal neighbor (with the same autonomous system number) at address 131. 108.234.2; and the third **neighbor** and **remote-as** commands specify a neighbor on a different autonomous system.

```

route-policy pass-all
pass
end-policy

```

```

router bgp 109
  address-family ipv4 unicast
    network
  131.
  108.0.0 255.
  0.0.0
    network 192.
  31.7.0 255.
  0.0.0
    neighbor
  131.
  108.200.1
    remote-as 167
    exit
  address-family ipv4 unicast
    route-policy pass-all in
    route-policy pass-out out
  neighbor
  131.
  108.234.2
    remote-as 109
    exit
  address-family ipv4 unicast
  neighbor
  150.
  136.64.19
    remote-as 99
    exit
  address-family ipv4 unicast
    route-policy pass-all in
    route-policy pass-all out

```

## BGP Confederation: Example

The following is a sample configuration that shows several peers in a confederation. The confederation consists of three internal autonomous systems with autonomous system numbers 6001, 6002, and 6003. To the BGP speakers outside the confederation, the confederation looks like a normal autonomous system with autonomous system number 666 (specified using the **bgp confederation identifier** command).

In a BGP speaker in autonomous system 6001, the **bgp confederation peers** command marks the peers from autonomous systems 6002 and 6003 as special eBGP peers. Hence, peers 171. 69.232.55 and 171. 69.232.56 get the local preference, next hop, and MED unmodified in the updates. The router at 160. 69.69.1 is a normal eBGP speaker, and the updates received by it from this peer are just like a normal eBGP update from a peer in autonomous system 666.

```

router bgp 6001
  bgp confederation identifier 666
  bgp confederation peers
    6002
    6003
  exit
  address-family ipv4 unicast
  neighbor 171.
  69.232.55
    remote-as 6002
    exit
  address-family ipv4 unicast
  neighbor 171.
  69.232.56
    remote-as 6003
    exit
  address-family ipv4 unicast
  neighbor
  160.
  69.69.1
    remote-as 777

```

In a BGP speaker in autonomous system 6002, the peers from autonomous systems 6001 and 6003 are configured as special eBGP peers. Peer 170. 70.70.1 is a normal iBGP peer, and peer 199.99.99.2 is a normal eBGP peer from autonomous system 700.

```
router bgp 6002
  bgp confederation identifier 666
  bgp confederation peers
    6001
    6003
  exit
  address-family ipv4 unicast
  neighbor
170.
70.70.1
  remote-as 6002
  exit
  address-family ipv4 unicast
  neighbor 171.
69.232.57
  remote-as 6001
  exit
  address-family ipv4 unicast
  neighbor 171.
69.232.56
  remote-as 6003
  exit
  address-family ipv4 unicast
  neighbor
199.
99.99.2
  remote-as 700
  exit
  address-family ipv4 unicast
  route-policy pass-all in
  route-policy pass-all out
```

In a BGP speaker in autonomous system 6003, the peers from autonomous systems 6001 and 6002 are configured as special eBGP peers. Peer 200. 200.200.200 is a normal eBGP peer from autonomous system 701.

```
router bgp 6003
  bgp confederation identifier 666
  bgp confederation peers
    6001
    6002
  exit
  address-family ipv4 unicast
  neighbor 171.
69.232.57
  remote-as 6001
  exit
  address-family ipv4 unicast
  neighbor 171.
69.232.55
  remote-as 6002
  exit
  address-family ipv4 unicast
  neighbor
200.
200.200.200
  remote-as 701
  exit
  address-family ipv4 unicast
  route-policy pass-all in
  route-policy pass-all out
```

The following is a part of the configuration from the BGP speaker 200.200.200.205 from autonomous system 701 in the same example. Neighbor 171.69.232.56 is configured as a normal eBGP speaker from autonomous system 666. The internal division of the autonomous system into multiple autonomous systems is not known to the peers external to the confederation.

```

router bgp 701
  address-family ipv4 unicast
  neighbor
171.
69.232.56
  remote-as 666
  exit
  address-family ipv4 unicast
  route-policy pass-all in
  route-policy pass-all out
  exit
  address-family ipv4 unicast
  neighbor
200.
200.200.205
  remote-as 701

```

## BGP Route Reflector: Example

The following example shows how to use an address family to configure internal BGP peer 10.1.1.1 as a route reflector client for both unicast and multicast prefixes:

```

router bgp 140
  address-family ipv4 unicast
  neighbor 10.1.1.1
  remote-as 140
  address-family ipv4 unicast
  route-reflector-client
  exit
  address-family ipv4 multicast
  route-reflector-client

```

## BGP MDT Address Family Configuration: Example

The following example shows how to configure an MDT address family in BGP:

```

router bgp 10

  bgp router-id 10.0.0.2
  address-family ipv4 unicast
  address-family vpv4 unicast
  address-family ipv4 mdt

  !
  neighbor 1.1.1.1

  remote-as 11
  update-source Loopback0
  address-family ipv4 unicast
  address-family vpv4 unicast
  address-family ipv4 md

```

!

## BGP Nonstop Routing Configuration: Example

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

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

## Best-External Path Advertisement Configuration: Example

The following example shows how to configure Best-External Path Advertisement:

```
router bgp 100
 address-family l2vpn vpls-vpws
  advertise best-external
end
```

## Primary Backup Path Installation: Example

The following example shows how to enable installation of primary backup path:

```
router bgp 100
 address-family l2vpn vpls-vpws
  additional-paths install backup
end
```

## Allocated Local Label Retention: Example

The following example shows how to retain the previously allocated local label for the primary path on the primary PE for 10 minutes after reconvergence:

```
router bgp 100
 address-family l2vpn vpls-vpws
  retain local-label 10
end
```

## Where to Go Next

For detailed information about BGP commands, see *Cisco IOS XR Routing Command Reference for the Cisco XR 12000 Series Router*

## Additional References

The following sections provide references related to implementing BGP for Cisco IOS XR software.

### Related Documents

Related Topic	Document Title
BGP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR Routing Command Reference for the Cisco XR 12000 Series Router</i>
Cisco Express Forwarding (CEF) commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco IOS XR IP Addresses and Services Command Reference for the Cisco XR 12000 Series Router</i>
MPLS VPN configuration information.	<i>Cisco IOS XR MPLS Configuration Guide for the Cisco XR 12000 Series Router</i>
Bidirectional Forwarding Detection (BFD)	<i>Cisco IOS XR Interface and Hardware Component Configuration Guide for the Cisco XR 12000 Series Router</i> and <i>Cisco IOS XR Interface and Hardware Component Command Reference for the Cisco XR 12000 Series Router</i>
Task ID information.	Configuring AAA Services on Cisco IOS XR Software module of <i>Cisco IOS XR System Security Configuration Guide for the Cisco XR 12000 Series Router</i>

### Standards

Standards	Title
draft-bonica-tcp-auth-05.txt	<i>Authentication for TCP-based Routing and Management Protocols</i> , by R. Bonica, B. Weis, S. Viswanathan, A. Lange, O. Wheeler
draft-ietf-idr-bgp4-26.txt	<i>A Border Gateway Protocol 4</i> , by Y. Rekhter, T.Li, S. Hares
draft-ietf-idr-bgp4-mib-15.txt	<i>Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4)</i> , by J. Hass and S. Hares
draft-ietf-idr-cease-subcode-05.txt	<i>Subcodes for BGP Cease Notification Message</i> , by Enke Chen, V. Gillet
draft-ietf-idr-avoid-transition-00.txt	<i>Avoid BGP Best Path Transitions from One External to Another</i> , by Enke Chen, Srihari Sangli

Standards	Title
draft-ietf-idr-as4bytes-12.txt	<i>BGP Support for Four-octet AS Number Space</i> , by Quaizar Vohra, Enke Chen
draft-nalawade-idr-mdt-safi-03.txt	<i>MDT SAFI</i> , by Gargi Nalawade and Arjun Sreekantiah

**MIBs**

MIBs	MIBs Link
—	To locate and download MIBs using Cisco IOS XR software, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: <a href="http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

**RFCs**

RFCs	Title
RFC 1700	<i>Assigned Numbers</i>
RFC 1997	BGP Communities Attribute
RFC 2385	Protection of BGP Sessions via the TCP MD5 Signature Option
RFC 2439	BGP Route Flap Damping
RFC 2545	Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
RFC 2796	BGP Route Reflection - An Alternative to Full Mesh IBGP
RFC 2858	Multiprotocol Extensions for BGP-4
RFC 2918	Route Refresh Capability for BGP-4
RFC 3065	Autonomous System Confederations for BGP
RFC 3392	Capabilities Advertisement with BGP-4
RFC 4271	<i>A Border Gateway Protocol 4 (BGP-4)</i>
RFC 4364	<i>BGP/MPLS IP Virtual Private Networks (VPNs)</i>
RFC 4724	<i>Graceful Restart Mechanism for BGP</i>

**Technical Assistance**

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>