



BGP Default Limits and Multiprotocol Address Handling

This chapter explains Border Gateway Protocol (BGP) default limits and multiprotocol address handling, such as advertising IPv4 routes over IPv6 networks. It also covers how to configure a Multicast Distribution Tree (MDT) address family session in BGP, with examples for default and non-default VRFs.

- [BGP default peer and prefix limits, on page 1](#)
- [IPv4 NLRI advertisement with IPv6 next hops in multiprotocol BGP networks, on page 3](#)
- [BGP MDT address family sessions, on page 19](#)

BGP default peer and prefix limits

BGP default limits are safeguards that

- restrict the maximum number of BGP peers and prefixes to prevent resource depletion,
- allow network operators to control neighbor and prefix scaling, and
- help maintain router stability in case of misconfiguration.

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
Support for increased number of BGP peers	Release 7.3.1	This feature is now enhanced to support 750 IPv4 and 750 IPv6 BGP peers.

BGP imposes maximum limits on how many neighbors you can configure on a router and on the number of prefixes the router accepts from each peer for each address family.

Default BGP peer and prefix limit values

These default limits apply to BGP configurations:

- Maximum number of BGP peers:
 - Default: 4000

- Configurable range: 1 to 15000

Use the **bgp maximum neighbor** command to change.

- Attempts to exceed the maximum or reduce it below the current number of peers will fail.

• *Table 2: Default maximum prefixes accepted per peer, by address family*

Address family	Default maximum prefixes
IPv4 unicast	1,048,576
IPv4 labeled-unicast	131,072
IPv4 tunnel	1,048,576
IPv6 unicast	524,288
IPv6 labeled-unicast	131,072
IPv4 multicast	131,072
IPv6 multicast	131,072
IPv4 MVPN	2,097,152
VPNv4 unicast	2,097,152
IPv4 MDT	131,072
VPNv6 unicast	1,048,576
L2VPN EVPN	2,097,152

You can override these values using the **maximum-prefix limit** command for each peer and address family.

How BGP default limits work

Summary

The key components involved in the process are:

- Router: Tracks peer and prefix counts.
- BGP neighbor: Sends prefix advertisements.
- Configuration commands: Allow you to adjust limits as needed.

BGP monitors the number of peers and received prefixes to enforce default or configured limits, protecting the router from excessive resource usage.

Workflow

These stages describe how BGP default limits work:

- The router checks the number of configured peers when changes are made.

- If a peer tries to exceed the maximum, the router rejects the configuration.
- For prefixes, the router monitors the count for each neighbor and address family.
- If the number of received prefixes exceeds the limit, either the default or the configured limit, the router sends a cease notification and terminates the BGP peering with that neighbor.
- If you lower the maximum-prefix limit after the session has started and the neighbor has already sent more prefixes than allowed, the router immediately sends a cease notification and ends the peering.

Result: BGP protects router resources and ensures reliable operation by enforcing these limits and terminating peerings that exceed them.

IPv4 NLRI advertisement with IPv6 next hops in multiprotocol BGP networks

Advertising IPv4 network layer reachability information (NLRI) with IPv6 next hops is a multiprotocol BGP feature that

- enables seamless communication between IPv4 devices across an IPv6 core
- supports both default and non-default virtual routing and forwarding (VRF) instances, and
- increases network flexibility and efficiency by allowing IPv4 routes to use IPv6 next hops without requiring address family alignment between network segments.

Table 3: Feature History Table

Feature Name	Release Information	Feature Description
Advertising IPv4 NLRI with IPv6 next hops in the non-default VRF	Release 24.4.1	Introduced in this release on: Fixed Systems(8700)(select variants only*) *Advertising IPv4 NLRI with IPv6 NextHops in the non-default VRF is now supported on Cisco 8712-MOD-M routers
Advertising IPv4 NLRI with IPv6 next hops in the non-default VRF	Release 24.2.11	This feature enhances network efficiency and security by allowing you to create default and non-default virtual routing tables. These tables isolate traffic through customized routing policies, allowing for the communication of IPv4 address family over IPv6 next hops specifically within non-default VRFs.

Advertising IPv4 NLRI with IPv6 next hops in MP-BGP networks	Release 7.3.1	With the capability of Multiprotocol BGP for exchanging IPv4 address family over IPv6 next hop, legacy IPv4 packets from edge devices can traverse through IPv6 core networks seamlessly. This feature allows IPv4 NLRI to be encoded and advertised over IPv6 next hop in a single BGP session. This feature supports exchanging the IPv4 address family over the IPv6 next hop only in the default VRF. Configure the ipv4 forwarding-enable command on the edge router connected to the core interface to encode IPv4 packets over IPv6 next hop.
--	---------------	--

You can use this feature in networks where devices with different address families coexist. For example, an IPv4 core may be surrounded by IPv6 devices, or the reverse. In these mixed environments, multiprotocol BGP (MP-BGP) allows you to advertise IPv4 routes using IPv6 next hops, so reachability information can cross address family boundaries. Common reasons for deploying such topologies include limited availability of IPv4 addresses for interface assignment and phased migration to an all-IPv6 network.

By enabling IPv4 forwarding on interfaces, you ensure that the type of address family used does not limit traffic flow across the network.

How IPv4 NLRI advertisement with IPv6 next hops works

Summary

The key components involved in the process are:

- BGP speakers: Routers that participate in multiprotocol BGP and advertise NLRI.
- VRFs: Routing tables that isolate traffic and allow for customized policy application.
- Next hop addresses: IPv6 addresses used as the next hop for IPv4 NLRI.

The router uses multiprotocol BGP to encode and advertise IPv4 routes with IPv6 next hops, supporting both default and non-default VRF instances. This allows legacy IPv4 packets to traverse an IPv6 network core.

Workflow

These stages describe how IPv4 NLRI advertisement with IPv6 next hops works:

- Configure interfaces with IPv6 addresses and enable IPv4 forwarding.
- Create route policies to set IPv6 next hops for IPv4 prefixes.
- Establish BGP neighbor relationships using IPv6 addresses.
- Advertise IPv4 NLRI with IPv6 next hop information through BGP sessions.
- Verify advertisement and routing through show commands.

Result: IPv4 routes are advertised and reachable over IPv6 next hops, supporting seamless traffic flow in mixed-protocol environments.

Configure IPv4 NLRI advertisement with IPv6 next hops in the default VRF

Enable advertisement of IPv4 NLRI with IPv6 next hops in the default VRF.

In the following example, the eBGP peer is configured on the HundredGigE 0/0/0/24 interface, and the iBGP peer is configured on the HundredGigE 0/0/0/25 interface of the router being configured.

Procedure

Step 1 Configure the required interfaces with IPv6 addresses and enable IPv4 forwarding.

Example:

```
Router# configure
Router(config)# interface HundredGigE 0/0/0/24
Router(config-if)# ipv6 address 2000::2/64
Router(config-if)# ipv4 forwarding-enable
Router(config-if)# no shutdown
Router(config-if)# commit
Router(config-if)# exit

Router(config)# interface HundredGigE 0/0/0/25
Router(config-if)# ipv6 address 3000::2/64
Router(config-if)# ipv4 forwarding-enable
Router(config-if)# no shutdown
Router(config-if)# commit
Router(config-if)# exit
```

Step 2 Create a prefix set and route policy to set IPv6 next hop for IPv4 prefixes.

Example:

```
Router(config)# route-policy 5549
Router(config-rpl)# if destination in 5549-pfx then set next-hop 20:20::20:200 endif
Router(config-rpl)# end-policy
Router(config)# prefix-set 5549-pfx
Router(config-pfx)# 3.3.3.3/32, 100.1.1.0/24 end-set
Router(config)# commit
```

Step 3 Configure iBGP and eBGP peers using IPv6 addresses and apply the route policies.

Example:

```
Router(config)# route-policy pass-all
Router(config-rpl)# pass
Router(config-rpl)# end-policy
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# exit
Router(config-bgp)# address-family ipv4 multicast
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 3000::1
Router(config-bgp-nbr)# remote-as 100 Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# next-hop-self Router(config-bgp-nbr-af)# route-policy pass-all in
Router(config-bgp-nbr-af)# route-policy pass-all out
Router(config-bgp-nbr-af)# commit
Router(config-bgp-nbr-af)# root
```

Step 4 Configure an eBGP peer through the HundredGigE 0/0/0/24 interface and use the route policy for setting an IPv6 nexthop for IPv4 routes.

Example:

```

Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# exit
Router(config-bgp)# address-family ipv4 multicast
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 2000::1
Router(config-bgp-nbr)# remote-as 200
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# route-policy pass-all in
Router(config-bgp-nbr-af)# route-policy 5549 out
Router(config-bgp-nbr-af)# commit
Router(config-bgp-nbr-af)# root

```

Step 5 Verify the running configuration.**Example:**

```

Router# show running-config
...
interface HundredGigE0/0/0/24 ipv4 forwarding-enable
ipv6 address 2000::2/64
!
interface HundredGigE0/0/0/25 ipv4 forwarding-enable
ipv6 address 3000::2/64
!
prefix-set 5549-pfx 3.3.3.3/32,
100.1.1.0/24
end-set
!
route-policy 5549
if destination in 5549-pfx then set next-hop 20:20::20:200
endif end-policy
!
route-policy pass-all pass
end-policy
!
router bgp 100
address-family ipv4 unicast
!
address-family ipv4 multicast
!
neighbor 2000::1
remote-as 200
address-family ipv4 unicast route-policy pass-all in route-policy 5549 out
!
!
neighbor 3000::1
remote-as 100
address-family ipv4 unicast route-policy pass-all in route-policy pass-all out next-hop-self
!
!
end

```

Step 6 Use these show commands to verify the BGP neighbor configuration and the advertisement of next hops.**Example:**

```

Router# show bgp neighbor
BGP neighbor is 10:10::10:10
Remote AS 100, local AS 100, internal link Remote router ID 10.10.10.10
Cluster ID 30.30.30.30

```

```

BGP state = Established, up for 11:42:17 NSR State: NSR Ready
Last read 00:00:10, Last read before reset 00:00:00 Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3 Last write 00:00:09, attempted
 19, written 19
Second last write 00:01:09, attempted 19, written 19
Last write before reset 00:00:00, attempted 0, written 0 Second last write before reset 00:00:00,
attempted 0, written 0
Last write pulse rcvd Jun 12 11:57:40.005 last full Jun 12 03:38:09.496 pulse count 1766

Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00 Last KA expiry before reset
00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00 Last KA start before reset 00:00:00, second
last 00:00:00 Precedence: internet
Non-stop routing is enabled
Multi-protocol capability received Neighbor capabilities:
Route refresh: advertised (old + new) and received (old + new) 4-byte AS: advertised and received
Address family IPv4 Unicast: advertised and received Received 866 messages, 0 notifications, 0 in
queue
Sent 1021 messages, 0 notifications, 0 in queue Minimum time between advertisement runs is 0 secs
Inbound message logging enabled, 3 messages buffered Outbound message logging enabled, 3 messages
buffered

For Address Family: IPv4 Unicast BGP neighbor version 120021
Update group: 0.2 Filter-group: 0.2 No Refresh request being processed Route-Reflector Client
Extended Nexthop Encoding: advertised and received
Route refresh request: received 0, sent 0
11 accepted prefixes, 11 are bestpaths Exact no. of prefixes denied : 0. Cumulative no. of prefixes
denied: 0.
Prefix advertised 60006, suppressed 0, withdrawn 60000 Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min AIGP is enabled
An EoR was not received during read-only mode
Last ack version 120021, Last synced ack version 120021 Outstanding version objects: current 0, max
 3 Additional-paths operation: None
Send Multicast Attributes
Advertise routes with local-label via Unicast SAFI

Connections established 1; dropped 0
Local host: 30:30::30:30, Local port: 51453, IF Handle: 0x00000000 Foreign host: 10:10::10:10, Foreign
port: 179
Last reset 00:00:00

```

```

Router# show bgp ipv4 unicast update-group
Mon Jun 12 11:47:31.543 UTC

```

```

Update group for IPv4 Unicast, index 0.2: Attributes:
Neighbor sessions are IPv6 Internal
Common admin
First neighbor AS: 100 Send communities
Send GSHUT community if originated Send extended communities
Route Reflector Client 4-byte AS capable
Advertise routes with local-label via Unicast SAFI Send AIGP
Send multicast attributes Extended Nexthop Encoding
Minimum advertisement interval: 0 secs Update group desynchronized: 0
Sub-groups merged: 5
Number of refresh subgroups: 0
Messages formatted: 156, replicated: 228 All neighbors are assigned to sub-group(s)
Neighbors in sub-group: 0.2, Filter-Groups num:1 Neighbors in filter-group: 0.2(RT num: 0)
10:10::10:10
20:20::20:20

```

```

Router# show bgp 3.3.3.3/32

```

Configure IPv4 NLRI advertisement with IPv6 next hops in the default VRF

```

Mon Jun 12 11:57:59.451 UTC
BGP routing table entry for 3.3.3.3/32 Versions:
Process bRIB/RIB SendTblVer
Speaker 21 21
Last Modified: Jun 12 00:15:45.314 for 11:42:14
Paths: (1 available, best #1)
Advertised to update-groups (with more than one peer): 0.2
Path #1: Received by speaker 0
Advertised to update-groups (with more than one peer): 0.2
3000, (Received from a RR-client)
20:20::20:20 (metric 1) from 20:20::20:20 (20.20.20.20)
Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best Received Path ID 0,
Local Path ID 0, version 21

Router# show bgp ipv4 unicast nexthops
...
..... Snippet .....

Gateway Address Family: IPv6 Unicast Table ID: 0xe0800000
Nexthop Count: 2
Critical Trigger Delay: 3000msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 3, RIB version: 1
EPE Table Version: 1, EPE Label version: 1
EPE Downloaded Version: 1, EPE Standby Version: 1

Status codes: R/UR Reachable/Unreachable
C/NC Connected/Not-connected L/NL Local/Non-local
PR Pending Registration I Invalid (Policy drop)
Next Hop Status Metric Tbl-ID Notf LastRIBEvent RefCount
10:10::10:10 [R][NC][NL] 1e0800000 1/0 11:42:43 (Cri) 11/14
20:20::20:20 [R][NC][NL] 1e0800000 1/0 11:42:43 (Cri) 8/11

Router# show bgp ipv4 unicast nexthops 10:10::10:10
Nexthop: 10:10::10:10 VRF: Default
Nexthop ID: 0x6000001, Version: 0x2 Nexthop Flags: 0x00000080
Nexthop Handle: 0x7f53e85b136c

RIB Related Information:
Firsthop interface handle 0x00000140
Gateway TBL Id: 0xe0800000 Gateway Flags: 0x00000080 Gateway Handle: 0x7f540d2e8f00
Gateway: reachable, non-Connected route, prefix length 128 Resolving Route: 10:10::10:10/128 (ospf
100)
Paths: 0
RIB Nexthop ID: 0x0
Status: [Reachable][Not Connected][Not Local] Metric: 1
Registration: Asynchronous, Completed: 2d21h Events: Critical (1)/Non-critical (0)
Last Received: 11:42:55 (Critical)
Last gw update: (Crit-notif) 11:42:55(rib) Reference Count: 11

Prefix Related Information Active Tables: [IPv4 Unicast] Metrics: [0x1]
Reference Counts: [11] Interface Handle: 0x0

Router# show route 3.3.3.3/32
Mon Jun 12 12:32:13.503 UTC

Routing entry for 3.3.3.3/32
Known via "bgp 100", distance 200, metric 0 Tag 3000, type internal
Installed Jun 12 00:15:45.626 for 12:16:28 Routing Descriptor Blocks
20:20::20:20, from 20:20::20:20
Route metric is 0 No advertising protos.

Router# show route 3.3.3.3/32 detail
Mon Jun 12 12:32:16.447 UTC

```



```

Routing entry for 3.3.3.3/32
Known via "bgp 100", distance 200, metric 0 Tag 3000, type internal
Installed Jun 12 00:15:45.628 for 12:16:30 Routing Descriptor Blocks
20:20::20:20, from 20:20::20:20
Route metric is 0 Label: None Tunnel ID: None
Binding Label: None
Extended communities count: 0 NHID:0x0(Ref:0)
Route version is 0x1 (1) No local label
IP Precedence: Not Set QoS Group ID: Not Set Flow-tag: Not Set
Fwd-class: Not Set
Route Priority: RIB_PRIORITY_RECURSIVE (12) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 4, Download Version 24 No advertising protos.

```

```

Router# show cef 3.3.3.3/32
Mon Jun 12 12:32:22.627 UTC
3.3.3.3/32, version 24, internal 0x5000001 0x0 (ptr 0x8e0b76b8) [1], 0x0 (0x0), 0x0 (0x0)
Updated Jun 12 00:15:45.631
local adjacency fe80::f83b:74ff:fe65:f004
Prefix Len 32, traffic index 0, precedence n/a, priority 4
via 20:20::20:20/128, 2 dependencies, recursive [flags 0x6000] path-idx 0 NHID 0x0 [0x8e352034 0x0]
next hop VRF - 'default', table - 0xe0800000 next hop 20:20::20:20/128 via 20:20::20:20/128

Router# show cef 3.3.3.3/32 detail
Mon Jun 12 12:32:25.415 UTC
3.3.3.3/32, version 24, internal 0x5000001 0x0 (ptr 0x8e0b76b8) [1], 0x0 (0x0), 0x0 (0x0)
Updated Jun 12 00:15:45.632
local adjacency fe80::f83b:74ff:fe65:f004
Prefix Len 32, traffic index 0, precedence n/a, priority 4
gateway array (0x8eec3170) reference count 6, flags 0x2010, source rib (7), 0 backups [1 type 3 flags
0x48501 (0x8e191998) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Jun 12 00:15:45.632
LDI Update time Jun 12 00:15:45.632
via 20:20::20:20/128, 2 dependencies, recursive [flags 0x6000] path-idx 0 NHID 0x0 [0x8e352034 0x0]
next hop VRF - 'default', table - 0xe0800000 next hop 20:20::20:20/128 via 20:20::20:20/128

Load distribution: 0 1 (refcount 1)

Hash OK Interface Address
0 Y FortyGigE0/0/0/0 fe80::f83b:74ff:fe65:f004
1 Y FortyGigE0/0/0/25 fe80::f83b:74ff:fe65:f08c

```

You have successfully configured and verified the advertisement of IPv4 NLRI through IPv6 nexthops in the default VRF.

Configure IPv4 NLRI advertisement with IPv6 next hops in a non-default VRF

Enable advertisement of IPv4 NLRI with IPv6 next hops in a non-default VRF.

Procedure

Step 1 Create and configure the VRF.

Example:

Configure IPv4 NLRI advertisement with IPv6 next hops in a non-default VRF

```
Router# configure
Router(config)# vrf RFC5549
Router(config-vrf)# address-family ipv4 unicast
Router(config-vrf-af)# exit
Router(config-vrf)# address-family ipv4 multicast
Router(config-vrf-af)# exit
Router(config-vrf)# commit
```

Step 2 Configure the required interfaces with an IPv6 address to the eBGP neighbor and the **ipv4 forwarding-enable** command.

Example:

```
Router# configure
Router(config)# interface GigabitEthernet0/1/0/0.2
Router(config-subif)# vrf RFC5549
Router(config-subif)# ipv4 forwarding-enable
Router(config-subif)# ipv6 address 2001:DB8:2::2/64
Router(config-subif)# encapsulation dot1q 2
Router(config-subif)# commit
```

Step 3 Configure the required interfaces with an IPv6 address to the iBGP neighbor and the **ipv4 forwarding-enable** command.

Example:

```
Router# configure
Router(config)# interface GigabitEthernet0/1/0/1.2
Router(config-subif)# vrf RFC5549
Router(config-subif)# ipv4 forwarding-enable
Router(config-subif)# ipv6 address 2001:DB9:2::2/64
Router(config-subif)# encapsulation dot1q 2
Router(config-subif)# commit
```

Step 4 Create a route policy to pass all.

Example:

```
Router# configure
Router(config)# route-policy pass-all
Router(config-rpl)# pass
Router(config-rpl)# end-policy
Router(config)# commit
```

Step 5 Configure BGP, VRF, and the required address families.

Example:

```
Router# configure
Router(config)# router bgp 2
Router(config-bgp)# bgp router-id 192.0.2.2
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# exit
Router(config-bgp)# address-family vpnv4 multicast
Router(config-bgp-af)# exit
Router(config-bgp)# vrf RFC5549
Router(config-bgp-vrf)# rd 1:100
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# exit
Router(config-bgp-vrf)# address-family ipv4 multicast
Router(config-bgp-vrf-af)# exit
Router(config-bgp-vrf)# commit
```

Step 6 Configure the eBGP neighbor.

Example:

```

Router# configure
Router(config)# router bgp 2
Router(config-bgp)# vrf RFC5549
Router(config-bgp-vrf)# neighbor 2001:DB8:2::1
Router(config-bgp-vrf-nbr)# remote-as 1
Router(config-bgp-vrf-nbr)# description eBGP neighbor in VRF
Router(config-bgp-vrf-nbr)# address-family ipv4 unicast
Router(config-bgp-vrf-nbr-af)# route-policy pass-all in
Router(config-bgp-vrf-nbr-af)# route-policy pass-all out
Router(config-bgp-vrf-nbr-af)# commit

```

Step 7 Configure the iBGP neighbor.

Example:

```

Router# configure
Router(config)# router bgp 2
Router(config-bgp)# vrf RFC5549
Router(config-bgp-vrf)# neighbor 2001:DB9:2::3
Router(config-bgp-vrf-nbr)# remote-as 2
Router(config-bgp-vrf-nbr)# description iBGP neighbor in VRF
Router(config-bgp-vrf-nbr)# address-family ipv4 unicast
Router(config-bgp-vrf-nbr-af)# route-policy pass-all in
Router(config-bgp-vrf-nbr-af)# route-policy pass-all out
Router(config-bgp-vrf-nbr-af)# commit

```

Verify IPv4 NLRI advertisement with IPv6 next hops in a non-default VRF

Verify advertisement of IPv4 NLRI with IPv6 next hops in a non-default VRF.

Procedure

Step 1 Verify the eBGP and iBGP VRF neighbors.

Example:

```

Router# show bgp vrf RFC5549 neighbors BGP neighbor is 2001:DB8:2::1, vrf RFC5549
Remote AS 1, local AS 2, external link Remote router ID 192.0.2.1
BGP state = Established, up for 00:03:37
Previous State: Idle
Last Received Message: KeepAlive NSR State: None
Last read 00:00:32, Last read before reset 00:05:07 Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3 Last write 00:00:32, attempted
19, written 19
Second last write 00:01:32, attempted 19, written 19
Last write before reset 00:05:06, attempted 19, written 19 Second last write before reset 00:06:06,
attempted 19, written 19
Last write pulse rcvd Mar 27 23:40:00.449 last full not set pulse count 79 Last write pulse rcvd
before reset 00:05:06
Last insert into reset queue: Mar 27 23:36:20.435, removed at Mar 27 23:36:20.435 Socket not armed
for io, armed for read, armed for write
Last write thread event before reset 00:04:12, second last 00:05:06 Last KA expiry before reset
00:05:06, second last 00:06:06
Last KA error before reset 00:00:00, KA not sent 00:00:00 Last KA start before reset 00:05:06, second
last 00:06:06 Precedence: internet
Non-stop routing is enabled Enforcing first AS is enabled Multi-protocol capability received Neighbor
capabilities:

```

Verify IPv4 NLRI advertisement with IPv6 next hops in a non-default VRF

Route refresh: advertised (old + new) and received (old + new) 4-byte AS: advertised and received
 Address family IPv4 Unicast: advertised and received Received 42 messages, 0 notifications, 0 in queue
 Sent 42 messages, 1 notifications, 0 in queue Minimum time between advertisement runs is 0 secs
 Inbound message logging enabled, 3 messages buffered Outbound message logging enabled, 3 messages buffered
 Fast fallover is enabled
 Neighbor is directly connected
 Neighbor fast-fallover is not configured
 Neighbor is external and fast-external-fallover is not disabled

For Address Family: IPv4 Unicast BGP neighbor version 10
 Update group: 0.3 Filter-group: 0.4 No Refresh request being processed AF-dependent capabilities:
Extended Nexthop Encoding: advertised and received
 Route refresh request: received 0, sent 0 Policy for incoming advertisements is pass-all Policy for outgoing advertisements is pass-all
1 accepted prefixes, 1 are bestpaths Exact no. of prefixes denied: 0 Cumulative no. of prefixes denied: 0
 Prefix advertised 1, suppressed 0, withdrawn 0 An EoR was received during read-only mode Last ack version 10, Last synced ack version 0
 Outstanding version objects: current 0, max 1, refresh 0 Additional-paths operation: None
 Advertise routes with local-label via Unicast SAFI Slow Peer State: Detection-only
 Detected state: FALSE, Detection threshold: 300 Detection Count: 0, Recovery Count: 0

Connections established 2; dropped 1

Local host: 2001:DB8:2::2, Local port: 18907, IF Handle: 0x008001e0 Foreign host: 2001:DB8:2::1, Foreign port: 179
 Last reset 00:04:12, due to Address family removed (CEASE notification sent - configuration change)
 Time since last notification sent to neighbor: 00:04:12 Error Code: configuration change
 Notification data sent:
 None

BGP neighbor is 2001:DB9:2::3, vrf RFC5549 Remote AS 2, local AS 2, internal link Remote router ID 192.0.2.3

BGP state = Established, up for 00:03:50

Previous State: Idle
 Last Received Message: KeepAlive NSR State: None
 Last read 00:00:45, Last read before reset 00:05:06 Hold time is 180, keepalive interval is 60 seconds
 Configured hold time: 180, keepalive: 60, min acceptable hold time: 3 Last write 00:00:32, attempted 19, written 19
 Second last write 00:01:32, attempted 19, written 19
 Last write before reset 00:05:06, attempted 19, written 19 Second last write before reset 00:06:06, attempted 19, written 19
 Last write pulse rcvd Mar 27 23:40:00.459 last full not set pulse count 80 Last write pulse rcvd before reset 00:05:06
 Last insert into reset queue: Mar 27 23:36:20.434, removed at Mar 27 23:36:20.434 Socket not armed for io, armed for read, armed for write
 Last write thread event before reset 00:05:06, second last 00:05:06 Last KA expiry before reset 00:05:06, second last 00:06:06
 Last KA error before reset 00:00:00, KA not sent 00:00:00 Last KA start before reset 00:05:06, second last 00:06:06 Precedence: internet
 Non-stop routing is enabled
 Multi-protocol capability received Neighbor capabilities:
 Route refresh: advertised (old + new) and received (old + new) 4-byte AS: advertised and received
 Address family IPv4 Unicast: advertised and received Received 42 messages, 0 notifications, 0 in queue
 Sent 44 messages, 1 notifications, 0 in queue Minimum time between advertisement runs is 0 secs
 Inbound message logging enabled, 3 messages buffered Outbound message logging enabled, 3 messages buffered
 Fast fallover is not enabled
 Neighbor is directly connected
 Neighbor fast-fallover is not configured

For Address Family: IPv4 Unicast BGP neighbor version 10
 Update group: 0.1 Filter-group: 0.3 No Refresh request being processed NEXT_HOP is always this router
 AF-dependent capabilities:

Extended Nexthop Encoding: advertised and received

Route refresh request: received 0, sent 0 Policy for incoming advertisements is pass-all Policy for outgoing advertisements is pass-all

1 accepted prefixes, 1 are bestpaths Exact no. of prefixes denied: 0 Cumulative no. of prefixes denied: 0

Prefix advertised 1, suppressed 0, withdrawn 0 An EoR was received during read-only mode Last ack version 10, Last synced ack version 0

Outstanding version objects: current 0, max 1, refresh 0 Additional-paths operation: None

Send Multicast Attributes

Advertise routes with local-label via Unicast SAFI Slow Peer State: Detection-only

Detected state: FALSE, Detection threshold: 300 Detection Count: 0, Recovery Count: 0

Connections established 2; dropped 1

Local host: 2001:DB9:2::2, Local port: 179, IF Handle: 0x00800240 Foreign host: 2001:DB9:2::3, Foreign port: 19682

Last reset 00:04:12, due to Address family removed (CEASE notification sent - configuration change)

Time since last notification sent to neighbor: 00:04:12 Error Code: configuration change

Notification data sent:

None

Step 2 Verify update groups for the eBGP and iBGP neighbors.**Example:**

Router# **show bgp vrf RFC5549 ipv4 unicast update-group**

Update group for IPv4 Unicast, index 0.1:

Attributes:

Neighbor sessions are IPv6 Outbound policy: pass-all Internal

Common admin

First neighbor AS: 2

Send communities

Send GSHUT community if originated Send extended communities

Next-hop-self enabled 4-byte AS capable

Advertise routes with local-label via Unicast SAFI Send multicast attributes

Extended Nexthop Encoding

Minimum advertisement interval: 0 secs Update group desynchronized: 0

Sub-groups merged: 0

Number of refresh subgroups: 0 Messages formatted: 2, replicated: 2

All neighbor are assigned to sub-group(s)

Neighbors in sub-group: 0.1, Filter-Groups num:1 Neighbors in filter-group: 0.3(RT num: 0)

2001:DB9:2::3

Update group for IPv4 Unicast, index 0.3: Attributes:

Outbound policy: pass-all First neighbor AS: 1 Directly connected IPv6 EBGP 4-byte AS capable

Advertise routes with local-label via Unicast SAFI

Extended Nexthop Encoding

Minimum advertisement interval: 0 secs Update group desynchronized: 0

Sub-groups merged: 0

Number of refresh subgroups: 0 Messages formatted: 2, replicated: 2

All neighbor are assigned to sub-group(s)

Neighbors in sub-group: 0.2, Filter-Groups num:1 Neighbors in filter-group: 0.4(RT num: 0)

2001:DB8:2::1

Step 3 Verify eBGP and iBGP IPv4 routes with IPv6 nexthops.**Example:**

Router# **show bgp vrf RFC5549 ipv4 unicast**

Wed Mar 27 23:41:44.494 IST

BGP VRF RFC5549, state: Active

BGP Route Distinguisher: 1:100

VRF ID: 0x60000003

BGP router identifier 100.1.1.2, local AS number 2

Non-stop routing is enabled BGP table state: Active

Table ID: 0xe0000012 RD version: 10 BGP table nexthop route policy:

Verify IPv4 NLRI advertisement with IPv6 next hops in a non-default VRF

```

BGP main routing table version 10
BGP NSR Initial initsync version 6 (Reached) BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, r RIB-failure, S stale, N Nexthop-discard Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path Route Distinguisher: 1:100 (default for vrf RFC5549)
Route Distinguisher Version: 10
*> 192.1.0.1/32 2001:DB8:2::1 0 0 1 ? N
*>i192.1.0.3/32 2001:DB9:2::3 0 100 0 ? N

Processed 2 prefixes, 2 paths

```

Step 4 Verify eBGP and iBGP routes.

a) Verify eBGP route.

Example:

```

Router# show bgp vrf RFC5549 ipv4 unicast 192.1.0.1/32
BGP routing table entry for 192.1.0.1/32, Route Distinguisher: 1:100
Versions:
Process bRIB/RIB SendTblVer
Speaker 10 10
Local Label: 24004
Last Modified: Mar 27 23:37:00.000 for 00:04:54 Last Delayed at: ---
Paths: (1 available, best #1)
Advertised to CE peers (in unique update groups):
2001:DB9:2::3
Path #1: Received by speaker 0
Advertised to CE peers (in unique update groups):
2001:DB9:2::3
2001:DB8:2::1 from 2001:DB8:2::1 (192.0.2.1)
Origin incomplete, metric 0, localpref 100, valid, external, best, group-best, import-candidate
Received Path ID 0, Local Path ID 1, version 10 Origin-AS validity: (disabled)

```

b) Verify iBGP route.

Example:

```

Router# show bgp vrf RFC5549 ipv4 unicast 192.1.0.3/32
BGP routing table entry for 192.1.0.3/32, Route Distinguisher: 1:100
Versions:
Process bRIB/RIB SendTblVer
Speaker 9 9
Local Label: 24007
Last Modified: Mar 27 23:36:47.000 for 00:05:29 Last Delayed at: ---
Paths: (1 available, best #1)
Advertised to CE peers (in unique update groups):
2001:DB8:2::1
Path #1: Received by speaker 0
Advertised to CE peers (in unique update groups): 2001:DB8:2::1
Local
2001:DB9:2::3 from 2001:DB9:2::3 (192.0.2.3)
Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, import-candidate
Received Path ID 0, Local Path ID 1, version 9

```

Step 5 Display IPv4 and IPv6 next-hop addresses for eBGP and iBGP IPv4 routes in the VRF.**Example:**

```

Router# show bgp vrf RFC5549 ipv4 unicast nexthops
Total Nexthop Processing Time Spent: 0.000 secs

Maximum Nexthop Processing Received: 00:00:00
Bestpaths Deleted: 0
Bestpaths Changed: 0 Time Spent: 0.000 secs

```

```

Last Notification Processing Received: 00:37:03
Time Spent: 0.000 secs IPv4 Unicast is active
Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000012
Gateway Reference Count: 1 Gateway AF Bits : 0x1 Nexthop Count: 0
Critical Trigger Delay: 50msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1
EPE Table Version: 1, EPE Label version: 0
EPE Downloaded Version: 0, EPE Standby Version: 0 IPv4 Unicast is active
Gateway Address Family: IPv6 Unicast
Table ID: 0xe0800012 Gateway Reference Count: 4 Gateway AF Bits : 0x1 Nexthop Count: 2
Critical Trigger Delay: 50msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 3, RIB version: 1
EPE Table Version: 1, EPE Label version: 0
EPE Downloaded Version: 0, EPE Standby Version: 0

Status codes: R/UR Reachable/Unreachable
C/NC Connected/Not-connected L/NL Local/Non-local

Next Hop PR I Pending Registration Invalid (Policy drop)
Status Metric

Tbl-ID

Notf

LastRIBEvent

RefCount
2001:DB8:2::1 [R][C][NL] 0 e0800012 1/0 00:37:03 (Cri) 1/3
2001:DB9:2::3 [R][C][NL] 0 e0800012 1/0 00:37:03 (Cri) 1/4
Counters
Next Hop Reachable Unreachable MetricIncrease MetricDecrease
2001:DB8:2::1 1 0 0 0
2001:DB9:2::3 1 0 0 0

```

Step 6 Verify IPv6 nexthop for the eBGP route.

Example:

```

Router# show bgp vrf RFC5549 ipv4 unicast nexthops 2001:DB8:2::1 Nexthop: 2001:DB8:2::1
VRF: RFC5549
Nexthop ID: 0x6000001, Version: 2 Nexthop Flags: 0x00020082 Nexthop Handle: 0x211b390
Tree Nexthop Handle: 0x211b390

Advertising neighbors:
2001:DB8:2::1

RIB Related Information:
Firsthop interface handle 0x008001e0
Gateway TBL Id: 0xe0800012 Gateway Flags: 0x00000080 Gateway Handle: 0x24c35a8
Gateway: reachable, Connected route, prefix length 64 Resolving Route: 2001:DB8:2::/64 (connected)
Paths: 1
RIB Nexthop ID: 0x20010 Nexthop sync slot: 23
Status: [Reachable][Connected][Not Local] Metric: 0
ORR afi bits: 0x0
Registration: Asynchronous, Completed: 00:37:16 Events: Critical (1)/Non-critical (0)

```

Verify IPv4 NLRI advertisement with IPv6 next hops in a non-default VRF

```
Last Received: 00:37:16 (Critical)
Last gw update: (Crit-notif) 00:37:16(rib)
Reference Count: 1
```

```
Reachable Notifications: 1 (last at Mar 27 23:06:16.748)
Unreachable Notifications: 0
Metric Increase Notifications: 0
Metric Decrease Notifications:
Most Recent Events: 0
Time Event Type Metric
Mar 27 23:06:16.748 Reachable 0
```

```
Prefix Related Information Active Tables: [IPv4 Unicast] Metrics: [0xffffffff] Reference Counts:
[1] Encapsulations: []
Interface Handle: 0x0
```

```
Linked Nexthop Count: 1 Attr ref-count: 3
```

Step 7 Verify IPv4 nexthop for the eBGP route.**Example:**

```
Router# show bgp vrf RFC5549 ipv4 unicast nexthops 2001:DB9:2::3
Nexthop: 2001:DB9:2::3
VRF: RFC5549
Nexthop ID: 0x6000002, Version: 3 Nexthop Flags: 0x00000082 Nexthop Handle: 0x211b130
Tree Nexthop Handle: 0x211b130

RIB Related Information:
Firsthop interface handle 0x00800240
Gateway TBL Id: 0xe0800012 Gateway Flags: 0x00000080 Gateway Handle: 0x24c34f0
Gateway: reachable, Connected route, prefix length 64 Resolving Route: 2001:DB9:2::/64 (connected)
Paths: 0
RIB Nexthop ID: 0x20011 Nexthop sync slot: 30
Status: [Reachable][Connected][Not Local] Metric: 0
ORR afi bits: 0x0
Registration: Asynchronous, Completed: 00:37:24 Events: Critical (1)/Non-critical (0)
Last Received: 00:37:24 (Critical)
Last gw update: (Crit-notif) 00:37:24(rib)
Reference Count: 1

Reachable Notifications: 1 (last at Mar 27 23:06:16.748)
Unreachable Notifications: 0
Metric Increase Notifications: 0
Metric Decrease Notifications:
Most Recent Events: 0
Time Event Type Metric
Mar 27 23:06:16.748 Reachable 0

Prefix Related Information Active Tables: [IPv4 Unicast] Metrics: [0xffffffff] Reference Counts:
[1] Encapsulations: []
Interface Handle: 0x0 Attr ref-count: 4
```

Step 8 Verify eBGP and iBGP routes in RIB.

- a) Display routing information for the specified prefix in the selected VRF.

Example:

```
Router# show route vrf RFC5549 192.1.0.1/32
Routing entry for 192.1.0.1/32
Known via "bgp 2", distance 20, metric 0
Tag 1
Local Label 24004, type external Installed Mar 27 23:37:00.244 for 00:09:00 Routing Descriptor
Blocks
```



```
fe80::83:9eff:fe38:c59f, from 2001:DB8:2::1, via GigabitEthernet0/1/0/0.2, BGP external Nexthop
in Vrf: "RFC5549", Table: "default", IPv6 Unicast, Table Id: 0xe0800012 Route metric is 0
No advertising protos.
```

```
Router# show route vrf RFC5549 192.1.0.3/32
Routing entry for 192.1.0.3/32
Known via "bgp 2", distance 200, metric 0
Local Label 24007, type internal Installed Mar 27 23:36:47.287 for 00:09:28 Routing Descriptor
Blocks
2001:DB9:2::3, from 2001:DB9:2::3
Nexthop in Vrf: "RFC5549", Table: "default", IPv6 Unicast, Table Id: 0xe0800012 Route metric is
0
No advertising protos.
```

- b) Display detailed routing information for the specified prefix in the selected VRF.

Example:

```
Router# show route vrf RFC5549 192.1.0.3/32 detail
Routing entry for 192.1.0.3/32
Known via "bgp 2", distance 20, metric 0
Tag 1, type external
Installed Mar 27 23:37:00.244 for 00:09:32 Routing Descriptor Blocks
fe80::83:9eff:fe38:c59f, from 2001:DB8:2::1, via GigabitEthernet0/1/0/0.2, BGP external Nexthop
in Vrf: "RFC5549", Table: "default", IPv6 Unicast, Table Id: 0xe0800012 Route metric is 0
Label: None Tunnel ID: None
Binding Label: None
Extended communities count: 0 NHID: 0x0 (Ref: 0)
Path Grouping ID: 1 Route version is 0x2 (2) Local Label: 0x5dc4 (24004) IP Precedence: Not Set
QoS Group ID: Not Set Flow-tag: Not Set Fwd-class: Not Set
Route Priority: RIB_PRIORITY_RECURSIVE (10) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 3, Download Version 59 No advertising protos.
```

```
Router# show route vrf RFC5549 192.1.0.3/32 detail
Routing entry for 192.1.0.3/32
Known via "bgp 2", distance 200, metric 0, type internal Installed Mar 27 23:36:47.287 for 00:09:52
Routing Descriptor Blocks
23:2:2::3, from 23:2:2::3
Nexthop in Vrf: "RFC5549", Table: "default", IPv6 Unicast, Table Id: 0xe0800012 Route metric is
0
Label: None Tunnel ID: None
Binding Label: None
Extended communities count: 0 NHID: 0x0 (Ref: 0)
Path Grouping ID: 2 Route version is 0x2 (2) Local Label: 0x5dc7 (24007) IP Precedence: Not Set
QoS Group ID: Not Set Flow-tag: Not Set Fwd-class: Not Set
Route Priority: RIB_PRIORITY_RECURSIVE (10) SVD Type RIB_SVD_TYPE_LOCAL
Download Priority 3, Download Version 57 No advertising protos.
```

Step 9 Verify eBGP and iBGP routes in Cisco Express Forwarding (CEF).

- a) Display CEF information for the specified prefix in the selected VRF.

Example:

```
Router# show cef vrf RFC5549 192.1.0.1/32
192.1.0.1/32, version 59, internal 0x1000001 0x30 (ptr 0x61c125a8) [1], 0x600 (0x6379bc58), 0xa20
(0x632aa3b8)
Updated Mar 27 23:37:00.247
remote adjacency to GigabitEthernet0/1/0/0.2
Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0x6441c190) reference count 2, flags 0x8068, source rib (7), 0 backups [3 type 4
flags 0x108401 (0x632c83b0) ext 0x0 (0x0)]
LW-LDI[type=1, refc=1, ptr=0x6379bc58, sh-ldi=0x632c83b0] gateway array update type-time 1 Mar 27
23:37:00.247
```

Verify IPv4 NLRI advertisement with IPv6 next hops in a non-default VRF

```

LDI Update time Mar 27 23:37:00.247
LW-LDI-TS Mar 27 23:37:00.247
via fe80::83:9eff:fe38:c59f/128, GigabitEthernet0/1/0/0.2, 4 dependencies, weight 0, class 0,
bgp-ext [flags 0x6020]
path-idx 0 NHID 0x0 [0x632e75a0 0x0]
next hop VRF - 'RFC5549', table - 0xe0800012 next hop fe80::83:9eff:fe38:c59f/128
remote adjacency
local label 24004 labels imposed {None} Load distribution: 0 (refcount 3)
Hash OK Interface Address
0 Y GigabitEthernet0/1/0/0.2 remote

Router# show cef vrf RFC5549 192.1.0.3/32
192.1.0.3/32, version 57, internal 0x1000001 0x30 (ptr 0x61c12698) [1], 0x0 (0x0), 0x200
(0x632aa1d8)
Updated Mar 27 23:36:47.289
remote adjacency to GigabitEthernet0/1/0/1.2
Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0x6441c098) reference count 2, flags 0xa078, source rib (7), 0 backups [1 type 4
flags 0x148501 (0x632c8358) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Mar 27 23:36:47.289
LDI Update time Mar 27 23:36:47.367
via 2001:DB9:2::3/128, 2 dependencies, recursive [flags 0x6000]
path-idx 0 NHID 0x0 [0x61c7a398 0x0]
next hop VRF - 'RFC5549', table - 0xe0800012 next hop 2001:DB9:2::3/128 via 2001:DB9:2::3/128
local label 24007
next hop 2001:DB9:2::3/128 Gi0/1/0/1.2 labels imposed {None}

Load distribution: 0 (refcount 1)
Hash OK Interface Address
0 Y GigabitEthernet0/1/0/1.2 remote

```

- b) Display detailed CEF information for the specified prefix in the selected VRF.

Example:

```

Router# show cef vrf RFC5549 192.1.0.1/32 detail
192.1.0.1/32, version 59, internal 0x1000001 0x30 (ptr 0x61c125a8) [1], 0x600 (0x6379bc58), 0xa20
(0x632aa3b8)
Updated Mar 27 23:37:00.247
remote adjacency to GigabitEthernet0/1/0/0.2
Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0x6441c190) reference count 2, flags 0x8068, source rib (7), 0 backups [3 type 4
flags 0x108401 (0x632c83b0) ext 0x0 (0x0)]
LW-LDI[type=1, refc=1, ptr=0x6379bc58, sh-ldi=0x632c83b0] gateway array update type-time 1 Mar 27
23:37:00.247
LDI Update time Mar 27 23:37:00.247
LW-LDI-TS Mar 27 23:37:00.247
via fe80::83:9eff:fe38:c59f/128, GigabitEthernet0/1/0/0.2, 4 dependencies, weight 0, class 0,
bgp-ext [flags 0x6020]
path-idx 0 NHID 0x0 [0x632e75a0 0x0]
next hop VRF - 'RFC5549', table - 0xe0800012 next hop fe80::83:9eff:fe38:c59f/128
remote adjacency
local label 24004 labels imposed {None}

Load distribution: 0 (refcount 3)

Hash OK Interface Address
0 Y GigabitEthernet0/1/0/0.2 remote

Router# show cef vrf RFC5549 192.1.0.3/32 detail
192.1.0.3/32, version 57, internal 0x1000001 0x30 (ptr 0x61c12698) [1], 0x0 (0x0), 0x200
(0x632aa1d8)
Updated Mar 27 23:36:47.289
remote adjacency to GigabitEthernet0/1/0/1.2

```

```

Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0x6441c098) reference count 2, flags 0xa078, source rib (7), 0 backups [1 type 4
flags 0x148501 (0x632c8358) ext 0x0 (0x0)]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
gateway array update type-time 1 Mar 27 23:36:47.289
LDI Update time Mar 27 23:36:47.367
via 2001:DB9:2::3/128, 2 dependencies, recursive [flags 0x6000] path-idx 0 NHID 0x0 [0x61c7a398
0x0]
next hop VRF - 'RFC5549', table - 0xe0800012 next hop 2001:DB9:2::3/128 via 2001:DB9:2::3/128
local label 24007
next hop 2001:DB9:2::3/128 Gi0/1/0/1.2 labels imposed {None}

Load distribution: 0 (refcount 1)

Hash OK Interface Address
0 Y GigabitEthernet0/1/0/1.2 remote

```

BGP MDT address family sessions

A multicast distribution tree (MDT) address family session in BGP is a feature that

- enables you to exchange IPv4 multicast routing information using the MDT subaddress family identifier (SAFI)
- supports multicast virtual private networks (MVPN), including both IPv4 and IPv6 deployments, and
- allows you to redistribute multicast routing information across multiple protocols and VRFs.

This feature is essential for service providers and enterprises that deploy multicast VPNs over BGP-based core networks.

Configure a BGP MDT address family session

Set up a BGP session for IPv4 MDT, which can also support multicast VPNv4 and VPNv6 distribution.

Before you begin

- Ensure you have decided on the autonomous system number (ASN) and router ID.
- Identify the interface or loopback address to use as the update source.

Procedure

Step 1 Enter BGP configuration mode for your autonomous system.

Example:

```

Router# configure
Router(config)# router bgp 120

```

- Step 2** Enable the MDT address family. If you are configuring multicast MVPN, also enable the unicast and VPN address families.

Example:

```
Router(config-vrf)# address-family ipv4 unicast
Router(config-bgp-af)# exit
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# exit
Router(config-bgp)# address-family ipv4 mdt
Router(config-bgp-af)# exit
```

- Step 3** Configure the BGP neighbor, specify its remote AS, and set the update source.

Example:

```
Router(config-bgp)# neighbor 172.168.40.24
Router(config-bgp-nbr)# remote-as 2002
Router(config-bgp-nbr)# update-source loopback 0
```

- Step 4** Activate the required address families for the neighbor.

Example:

```
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# exit
Router(config-bgp-nbr)# address-family vpnv4 unicast
Router(config-bgp-nbr-af)# exit
Router(config-bgp)# address-family ipv4 mdt
Router(config-bgp-af)# exit
```

- Step 5** If you are configuring multicast MVPN, create a VRF instance and configure a route distinguisher (RD).

Example:

```
Router(config-bgp)# vrf vpn1
Router(config-bgp-vrf)# rd 1:1
```

- Step 6** Within the VRF, redistribute routes from another routing protocol, such as EIGRP or OSPF, into BGP.

Example:

```
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# redistribute eigrp 23
```

- Step 7** Commit the configuration changes.

Example:

```
Router(config-bgp-vrf-af)#commit
```

- Step 8** Run the **show running-config** command to verify the configuration.

Example:

```
Router# show running-config
router bgp 120
  address-family ipv4 mdt
  !
  address-family vpnv4 unicast
  !
  vrf vpn1
    rd 1:1
    address-family ipv4 unicast
    redistribute eigrp 23
  !
!
```

```
neighbor 172.168.40.24
  remote-as 2002
  update-source Loopback0
  address-family ipv4 mdt
  !
  address-family vpnv4 unicast
  !
  address-family ipv4 unicast
  !
!
```
