



VXLAN EVPN and TRM with IPv6 Multicast Underlay


This chapter contains the following sections:

- [Information About Configuring VXLAN EVPN and TRM with IPv6 in the Multicast Underlay](#)
- [Guidelines and Limitations for VXLAN EVPN and TRM with IPv6 in the Multicast Underlay](#)
- [Configuring VXLAN EVPN and TRM with IPv6 in the Multicast Underlay](#)
- [Example Configuration for VXLAN EVPN and TRM with IPv6 in the Multicast Underlay](#)
- [Verifying VXLAN EVPN and TRM with IPv6 Multicast Underlay](#)

Information About Configuring VXLAN EVPN and TRM with IPv6 Multicast Underlay

Cisco NX-OS Release 10.3(99x)F supports VXLAN with IPv6 Multicast in the Underlay. Hosts in the overlay can be IPv4 or IPv6. This requires IPv6 versions of the unicast routing protocols and utilizing IPv6 multicast in the underlay (PIMv6). Any multi-destination overlay traffic (such as TRM, BUM) can use the IPv6 multicast underlay.

Figure 1. Topology - VXLAN EVPN with IPv6 Multicast Underlay



The above topology shows four leafs and two spines in a VXLAN EVPN fabric. The underlay is IPv6 Multicast running PIMv6. RP is positioned in the spine with anycast RP.

Guidelines and Limitations for VXLAN EVPN and TRM with IPv6 Multicast Underlay

VXLAN EVPN and TRM with IPv6 Multicast Underlay has the following guidelines and limitations:

- The upgrade from nxos64-cs.10.3.99w.F.bin to nxos64-cs.10.3.99x.F.bin will be a disruptive upgrade.
- Spine-based static RP is supported in underlay.
- Cisco Nexus 9300-FX, FX2 and FX3 switches are supported as the leaf VTEP.
- N9K-X9716D-GX and N9K-X9736C-FX line cards are supported on the spine (EoR) and N9K-C9332D-GX2B is supported on the super spine.

Note: EoR requires configuring a non-default template using one of the following commands in global configuration mode:

- **system routing template-multicast-heavy**
- **system routing template-multicast-ext-heavy**
- OSPFv3, eBGP underlay is supported.
- PIMv6 ASM (sparse mode) is supported in underlay.
- PIMv6 Anycast RP is supported in underlay as RP redundancy.
- Underlay IPv6 Multicast is supported.
- For overlay traffic, each Cisco Nexus 9000 leaf is an RP. External RP is also supported.

Configuring VXLAN EVPN and TRM with IPv6 Multicast Underlay

Configuring IPv6 multicast underlay in the VXLAN fabric involves the following:

- Configuring L2-VNI based multicast group in underlay.
- Configuring L3-VNI based multicast group in underlay.
- Enabling PIMv6 for underlay.

Configuring L2-VNI based multicast group in underlay:

Under NVE configuration on a leaf, IPv6 multicast group (IPv6) is configured for each L2-VNI (VLAN).

SUMMARY STEPS

1. `configure terminal`
2. `interface nve1`
3. `member vni <vni>`
4. `mcast-group ipv6-prefix`
5. `exit`

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code> Example: <code>switch# configure terminal</code>	Enter configuration mode.
Step 2	<code>interface nve1</code> Example: <code>switch(config)# interface nve 1</code>	Configure the NVE interface.
Step 3	<code>member vni <vni></code> Example: <code>switch(config-if-nve)# member vni 10501</code>	Configure the Layer 2 virtual network identifier.
Step 4	<code>mcast-group ipv6-prefix</code> Example: <code>switch(config-if-nve-vni)# mcast-group ff10::1</code>	Builds the default multicast distribution tree for the Layer 2 VNI.
Step 5	<code>exit</code>	Exits the command mode.

	Command or Action	Purpose
	Example: switch(config-if-nve-vni) # exit	

Configuring L3-VNI based multicast group in underlay:

IPv6 multicast group (IPv6) is configured for each L3-VNI (VRF).

SUMMARY STEPS

1. configure terminal
2. interface nve1
3. member vni <vni> associate-vrf
4. mcast-group <ipv6-prefix>
5. exit

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal	Enter configuration mode.
Step 2	interface nve1 Example: switch(config) # interface nve 1	Configure the NVE interface.
Step 3	member vni vni associate-vrf Example: switch(config-if-nve) # member vni 50001 associate-vrf	Associates L3VNI to VRF.
Step 4	mcast-group ipv6-prefix Example: switch(config-if-nve-vni) # mcast-group ff10:0:0:1::1	Builds the default multicast distribution tree for the VRF VNI (Layer 3 VNI used for TRM).
Step 5	exit Example: switch(config-if-nve-vni) # exit	Exits command mode.

Enabling PIMv6 for underlay:

PIMv6 in underlay is configured as follows:

SUMMARY STEPS

1. configure terminal
2. interface loopback <number>
3. ipv6 address <ipv6-prefix>

4. `ipv6 pim sparse-mode`
5. `interface nve1`
6. `source-interface loopback <number>`
7. `exit`

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal	Enters global configuration mode.
Step 2	interface loopback <i>number</i> Example: switch(config)# interface loopback 1	Configures an interface loopback. This example configures interface loopback 1.
Step 3	ipv6 address <i>ipv6-prefix</i> Example: switch(config-if)# ipv6 address 11:0:0:1::1/128	Configures an IP address for this interface. It should be a unique IP address that helps to identify this router.
Step 4	ipv6 pim sparse-mode Example: switch(config-if)# ipv6 pim sparse-mode	Enable PIM6 sparse mode.
Step 5	interface nve1 Example: switch(config-if-nve)# interface nve 1	Configure the NVE interface.
Step 6	source-interface loopback <i>number</i> Example: switch(config-if-nve)# source-interface loopback 1	Configures an source interface loopback.
Step 7	exit Example: switch(config-if-nve)# exit	Exits command mode.

Note:

For the PIMv6 configuration see the [Cisco Nexus 9000 Series NX-OS Multicast Routing Configuration Guide, Release 10.3\(x\)](#).

For the TRM configuration see the [Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide, Release 10.3\(x\)](#).

Example Configuration for VXLAN EVPN and TRM with IPv6 Multicast Underlay

In the following section, the sample configuration for the leaf, spine and RP are shown.

Leaf - sample configuration of IPv6 multicast underlay:

NVE Configuration

```
interface nve1
  no shutdown
  host-reachability protocol bgp
  source-interface loopback1
  member vni 10501
    mcast-group ff10::1
  member vni 50001 associate-vrf
    mcast-group ff10:0:0:1::1
```

PIMv6 Configuration

```
feature pim6

ipv6 pim rp-address 101:101:101:101::101 group-list ff00::/8

interface loopback1
  ipv6 address 172:172:16:1::1/128
  ipv6 pim sparse-mode

interface Ethernet1/27
  ipv6 address 27:50:1:1::1/64
  ospfv3 hello-interval 1
  ipv6 router ospfv3 v6u area 0.0.0.0
  ipv6 pim sparse-mode
  no shutdown
```

BGP Configuration

```
router bgp 100
  router-id 172.16.1.1
  address-family ipv4 unicast
    maximum-paths 64
    maximum-paths ibgp 64
  address-family ipv6 unicast
    maximum-paths 64
    maximum-paths ibgp 64
  address-family ipv4 mvpn
  address-family l2vpn evpn
  neighbor 172:17:1:1::1
    remote-as 100
    update-source loopback0
  address-family ipv4 mvpn
    send-community
    send-community extended
  address-family ipv6 mvpn
    send-community
    send-community extended
  address-family l2vpn evpn
    send-community
```

```

neighbor 172:17:2:2::1
  remote-as 100
  update-source loopback0
  address-family ipv4 mvpn
    send-community
    send-community extended
  address-family ipv6 mvpn
    send-community
    send-community extended
  address-family l2vpn evpn
    send-community
    send-community extended
vrf VRF1
  reconnect-interval 1
  address-family ipv4 unicast
    network 150.1.1.1/32
    advertise l2vpn evpn
    redistribute hmm route-map hmmAdv

evpn
  vni 10501 l2
  rd auto
  route-target import auto
  route-target export auto
vrf context VRF1
  vni 50001
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto mvpn
    route-target both auto evpn
  address-family ipv6 unicast
    route-target both auto
    route-target both auto mvpn
    route-target both auto evpn

```

Spine - sample configuration of IPv6 multicast underlay:

NVE Configuration

```
nv overlay evpn
```

PIMv6 Configuration

```

feature pim6

ipv6 pim rp-address 101:101:101:101::101 group-list ff00::/8
ipv6 pim anycast-rp 101:101:101:101::101 102:102:102:102::102
ipv6 pim anycast-rp 101:101:101:101::101 103:103:103:103::103

interface loopback101
  ipv6 address 101:101:101:101::101/128
  ipv6 router ospfv3 v6u area 0.0.0.0
  ipv6 pim sparse-mode

interface loopback102
  ipv6 address 102:102:102:102::102/128
  ipv6 router ospfv3 v6u area 0.0.0.0
  ipv6 pim sparse-mode

```

```
interface Ethernet1/50/1
  ipv6 address 27:50:1:1::2/64
  ipv6 pim sparse-mode
  no shutdown
```

BGP Configuration

```
feature bgp

router bgp 100
  router-id 172.16.40.1
  address-family ipv4 mvpn
  address-family ipv6 mvpn
  address-family l2vpn evpn
  neighbor 172:16:1:1::1
    remote-as 100
    update-source loopback0
  address-family ipv4 mvpn
    send-community
    send-community extended
    route-reflector-client
  address-family ipv6 mvpn
    send-community
    send-community extended
    route-reflector-client
  address-family l2vpn evpn
    send-community
    send-community extended
    route-reflector-client
```

Verifying VXLAN EVPN and TRM with IPv6 Multicast Underlay

The following example is used to verify the status of the IPv6 Multicast Underlay configuration.

```
switch(config)# show run interface nve 1

!Command: show running-config interface nve1
!Running configuration last done at: Wed Jul  5 10:03:58 2023
!Time: Wed Jul  5 10:04:01 2023
version 10.3(99x) Bios:version 01.08

interface nve1
  no shutdown
  host-reachability protocol bgp
  source-interface loopback1
  member vni 10501
    mcast-group ff10::1
  member vni 50001 associate-vrf
    mcast-group ff10:0:0:1::1
```

Use the following command for verifying PIMv6 ASM configuration:

```
switch(config)# show ipv6 mroute
IPv6 Multicast Routing Table for VRF "default"

(*, ff10::1/128), uptime: 05:20:19, nve pim6 ipv6
  Incoming interface: Ethernet1/36, RPF nbr: fe80::23a:9cff:fe23:8367
  Outgoing interface list: (count: 1)
    nve1, uptime: 05:20:19, nve
```



```
(172:172:16:1::1/128, ff10::1/128), uptime: 05:20:19, nve m6rib pim6 ipv6
  Incoming interface: loopback1, RPF nbr: 172:172:16:1::1
  Outgoing interface list: (count: 2)
    Ethernet1/36, uptime: 01:47:03, pim6
    Ethernet1/27, uptime: 04:14:20, pim6

(*, ff10:0:0:1::10/128), uptime: 05:20:18, nve ipv6 pim6
  Incoming interface: Ethernet1/36, RPF nbr: fe80::23a:9cff:fe23:8367
  Outgoing interface list: (count: 1)
    nve1, uptime: 05:20:18, nve

(172:172:16:1::1/128, ff10:0:0:1::10/128), uptime: 05:20:18, nve m6rib ipv6
pim6
  Incoming interface: loopback1, RPF nbr: 172:172:16:1::1
  Outgoing interface list: (count: 2)
    Ethernet1/36, uptime: 04:04:35, pim6
    Ethernet1/27, uptime: 04:13:35, pim6
```

```
switch(config)# show ipv6 pim neighbor
PIM Neighbor Status for VRF "default"
Neighbor                               Interface      Uptime      Expires      DR
Bidir- BFD      ECMP Redirect
                                           Priority

Capable State  Capable
fe80::23a:9cff:fe28:5e07  Ethernet1/27  20:23:38  00:01:44  1
yes      n/a      no
  Secondary addresses:
    27:50:1:1::2
```

```
switch(config)# show ipv6 pim rp
PIM RP Status Information for VRF "default"
BSR disabled
BSR RP Candidate policy: None
BSR RP policy: None

RP: 101:101:101:101::101, (0),
  uptime: 21:30:43  priority: 255,
  RP-source: (local),
  group ranges:
    ff00::/8
```

The following example provides the output for leaf BGP neighbor-1.

```
switch(config-if)# show ipv6 bgp neighbors

BGP neighbor is 33:52:1:1::2, remote AS 200, ebgp link, Peer index 3
  BGP version 4, remote router ID 172.17.1.1
  Neighbor previous state = OpenConfirm
  BGP state = Established, up for 00:00:16
  Neighbor vrf: default
  Peer is directly attached, interface Ethernet1/33
  Enable logging neighbor events
  Last read 0.926823, hold time = 3, keepalive interval is 1 seconds
  Last written 0.926319, keepalive timer expiry due 0.073338
  Received 23 messages, 0 notifications, 0 bytes in queue
```

Sent 67 messages, 0 notifications, 0(0) bytes in queue
Enhanced error processing: On
0 discarded attributes
Connections established 1, dropped 0
Last update recd 00:00:15, Last update sent = 00:00:15
Last reset by us 00:08:45, due to session closed
Last error length sent: 0
Reset error value sent: 0
Reset error sent major: 104 minor: 0
Notification data sent:
Last reset by peer never, due to No error
Last error length received: 0
Reset error value received 0
Reset error received major: 0 minor: 0
Notification data received:

Neighbor capabilities:

Dynamic capability: advertised (mp, refresh, gr) received (mp, refresh, gr)
Dynamic capability (old): advertised received
Route refresh capability (new): advertised received
Route refresh capability (old): advertised received
4-Byte AS capability: advertised received
Address family IPv6 Unicast: advertised received
Graceful Restart capability: advertised received

Graceful Restart Parameters:

Address families advertised to peer:
IPv6 Unicast
Address families received from peer:
IPv6 Unicast
Forwarding state preserved by peer for:
Restart time advertised to peer: 400 seconds
Stale time for routes advertised by peer: 300 seconds
Restart time advertised by peer: 120 seconds
Extended Next Hop Encoding Capability: advertised received
Receive IPv6 next hop encoding Capability for AF:
IPv4 Unicast VPNv4 Unicast

Message statistics:

	Sent	Rcvd
Opens:	46	1
Notifications:	0	0
Updates:	2	2
Keepalives:	18	18
Route Refresh:	0	0
Capability:	2	2
Total:	67	23
Total bytes:	521	538
Bytes in queue:	0	0

For address family: IPv6 Unicast

BGP table version 10, neighbor version 10
3 accepted prefixes (3 paths), consuming 864 bytes of memory
0 received prefixes treated as withdrawn
2 sent prefixes (2 paths)
Inbound soft reconfiguration allowed(always)
Allow my ASN 3 times
Last End-of-RIB received 00:00:01 after session start
Last End-of-RIB sent 00:00:01 after session start
First convergence 00:00:01 after session start with 2 routes sent

```
Local host: 33:52:1:1::1, Local port: 179
Foreign host: 33:52:1:1::2, Foreign port: 17226
fd = 112
```

The following example provides the output for leaf BGP neighbor-2.

```
switch(config-if)# show bgp l2vpn evpn neighbors 172:17:1:1::1

BGP neighbor is 172:17:1:1::1, remote AS 200, ebgp link, Peer index 5
  BGP version 4, remote router ID 172.17.1.1
  Neighbor previous state = OpenConfirm
  BGP state = Established, up for 00:01:33
  Neighbor vrf: default
  Using loopback0 as update source for this peer
  Using iod 65 (loopback0) as update source
  Enable logging neighbor events
  External BGP peer might be up to 5 hops away
  Last read 0.933565, hold time = 3, keepalive interval is 1 seconds
  Last written 0.915927, keepalive timer expiry due 0.083742
  Received 105 messages, 0 notifications, 0 bytes in queue
  Sent 105 messages, 0 notifications, 0(0) bytes in queue
  Enhanced error processing: On
    0 discarded attributes
  Connections established 1, dropped 0
  Last update recd 00:01:32, Last update sent = 00:01:32
  Last reset by us never, due to No error
  Last error length sent: 0
  Reset error value sent: 0
  Reset error sent major: 0 minor: 0
  Notification data sent:
  Last reset by peer never, due to No error
  Last error length received: 0
  Reset error value received 0
  Reset error received major: 0 minor: 0
  Notification data received:

Neighbor capabilities:
  Dynamic capability: advertised (mp, refresh, gr) received (mp, refresh, gr)
  Dynamic capability (old): advertised received
  Route refresh capability (new): advertised received
  Route refresh capability (old): advertised received
  4-Byte AS capability: advertised received
  Address family IPv4 MVPN: advertised received
  Address family IPv6 MVPN: advertised received
  Address family L2VPN EVPN: advertised received
  Graceful Restart capability: advertised received

Graceful Restart Parameters:
  Address families advertised to peer:
    IPv4 MVPN  IPv6 MVPN  L2VPN EVPN
  Address families received from peer:
    IPv4 MVPN  IPv6 MVPN  L2VPN EVPN
  Forwarding state preserved by peer for:
  Restart time advertised to peer: 400 seconds
  Stale time for routes advertised by peer: 300 seconds
  Restart time advertised by peer: 120 seconds
  Extended Next Hop Encoding Capability: advertised received
  Receive IPv6 next hop encoding Capability for AF:
    IPv4 Unicast  VPNv4 Unicast
```

Message statistics:

	Sent	Rcvd
Opens:	1	1
Notifications:	0	0
Updates:	6	3
Keepalives:	95	95
Route Refresh:	0	0
Capability:	6	6
Total:	105	105
Total bytes:	2551	2047
Bytes in queue:	0	0

For address family: IPv4 MVPN

BGP table version 3, neighbor version 3
0 accepted prefixes (0 paths), consuming 0 bytes of memory
0 received prefixes treated as withdrawn
0 sent prefixes (0 paths)
Community attribute sent to this neighbor
Extended community attribute sent to this neighbor
Allow my ASN 3 times
Outbound route-map configured is RN_NextHop_Unchanged, handle obtained
Last End-of-RIB received 00:00:01 after session start
Last End-of-RIB sent 00:00:01 after session start
First convergence 00:00:01 after session start with 0 routes sent

For address family: IPv6 MVPN

BGP table version 3, neighbor version 3
0 accepted prefixes (0 paths), consuming 0 bytes of memory
0 received prefixes treated as withdrawn
0 sent prefixes (0 paths)
Community attribute sent to this neighbor
Extended community attribute sent to this neighbor
Allow my ASN 3 times
Outbound route-map configured is RN_NextHop_Unchanged, handle obtained
Last End-of-RIB received 00:00:01 after session start
Last End-of-RIB sent 00:00:01 after session start
First convergence 00:00:01 after session start with 0 routes sent

For address family: L2VPN EVPN

BGP table version 7, neighbor version 7
0 accepted prefixes (0 paths), consuming 0 bytes of memory
0 received prefixes treated as withdrawn
4 sent prefixes (4 paths)
Community attribute sent to this neighbor
Extended community attribute sent to this neighbor
Allow my ASN 3 times
Advertise GW IP is enabled
Outbound route-map configured is RN_NextHop_Unchanged, handle obtained
Last End-of-RIB received 00:00:01 after session start
Last End-of-RIB sent 00:00:01 after session start
First convergence 00:00:01 after session start with 4 routes sent

Local host: 172:16:1:2::1, Local port: 21132
Foreign host: 172:17:1:1::1, Foreign port: 179
fd = 113