



# Configure VXLAN with IPv6 in the Underlay (VXLANv6)

---

This chapter contains these sections:

- [VXLANv6 deployments, on page 1](#)
- [VXLANv6 vMACs, on page 2](#)
- [VXLANv6 vPC peer keepalives, on page 3](#)
- [VXLANv6 deployment features and limitations, on page 3](#)
- [Configure the VTEP IP address, on page 6](#)
- [Configure vPC for VXLANv6, on page 7](#)
- [VXLANv6 configuration requirements and examples, on page 9](#)
- [Verification commands for VXLANv6, on page 10](#)

## VXLANv6 deployments

A VXLAN with IPv6 in the Underlay (VXLANv6) deployment is a network virtualization solution that

- uses IPv6 in the underlay network and for VTEP addresses,
- supports overlay hosts that can be either IPv4 or IPv6, and
- enables modern routing protocols and BGP unnumbered peering in the underlay.

VXLAN BGP EVPN is deployed with IPv4 underlay and IPv4 VTEP. Hosts in the overlay can be IPv4 or IPv6. Support is added for VXLANv6 with an IPv6 VTEP. This requires IPv6 versions of the unicast routing protocols and utilizing ingress replication or multicast underlay for multi-destination traffic (BUM) in the underlay.

This solution is targeted for deployments where the VTEP is IPv6 only and the underlay is IPv6. The BGP sessions between the leaf and spine are also IPv6. The overlay hosts can be either IPv4 or IPv6.

VXLANv6 feature supports BGP unnumbered peering in the underlay.

### Supported protocols

These protocols are supported in the underlay:

- IS-IS

- OSPFv3
- eBGP

## VXLANv6 vMACs

VXLANv6 vMACs are virtual MAC address mechanisms that

- uniquely identify virtual tunnel endpoints (VTEPs) within a VXLANv6 fabric,
- are normally autogenerated from the configured VIP address, and
- support manual configuration for interoperability or uniqueness.

vPC VTEPs use vMAC (virtual MAC) with the VIP/PIP feature. vMAC is used with VIP and the system MAC is used with PIP.

### Key characteristics

- In IPv4 underlay, the vMAC is derived from the IPv4 VIP address using the formula:  

$$\text{VMAC} = 0x02 + 4 \text{ bytes IPv4 VIP address.}$$
- In IPv6 underlay, because the VIP is 128 bits and a MAC is only 48 bits, the system auto-generates the vMAC:  

$$\text{Autogenerated vMAC} = 0x06 + \text{the last 4 bytes of the IPv6 VIP address.}$$

### Interoperability

- If two vPC complexes have different VIPs but the same last 4 bytes of their IPv6 VIP, they will generate the same vMAC. This can result in vMAC "flapping" when a remote VTEP observes changes between the two VIPs.
- Cisco Nexus 9000 Series switches that support VXLAN IPv6 can handle such vMAC flapping. For interoperability with other vendor switches, administrators can manually configure vMACs to avoid conflicts.
- Manual configuration is performed using the `virtual-rmac <48-bit mac address>` command in the NVE interface configuration.
- The default behavior for VXLANv6 is to autogenerate and advertise the vMAC based on the VIP. Manual configuration overrides this default.

Example configuration:

```
interface nve1
    virtual-rmac <48-bit mac address>
```

- If the IPv4 VIP is 10.1.2.3, the vMAC will be 0x02 + 10.1.2.3 (in bytes).
- If the IPv6 VIP ends in ::1:2:3:4, the vMAC will be 0x06 + 0.0.3.4.

Two complexes with VIPs ending in the same last 4 bytes (for example, ::1:2:3:4 and ::5:6:3:4) generate the same vMAC, which may cause operational problems with some non-Cisco hardware.

# VXLANv6 vPC peer keepalives

VXLANv6 vPC peer keepalive is a network communication mechanism that

- monitors the health of both switches in a vPC pair,
- allows configuration using either IPv4 or IPv6 addresses on management or other network interfaces, and
- ensures reliable operation by requiring both peers to be configured with reachable, global unicast addresses.

## Key characteristics

- The peer-keepalive link operates only when both peers have correct and reachable addresses (either IPv4 or IPv6).
- This mechanism can function over in-band or out-of-band interfaces.
- The address used for peer-keepalive must always be a global unicast address.

To configure vPC peer keepalive using IPv6 addresses, use this command:

```
vpc domain 1
peer-keepalive destination 2001:DB8::1 source 2001:DB8::2 vrf management
```

# VXLANv6 deployment features and limitations

VXLAN with IPv6 in the underlay (VXLANv6) introduces specific guidelines, platform support, feature availability, and limitations. Use this reference to understand deployment boundaries, supported and unsupported features, and hardware compatibility.

## Deployment and configuration guidelines

Key deployment and configuration limitations for VXLANv6 include:

- Dual Stack (IPv4 and IPv6) is not supported for VXLAN underlay. It should either be IPv4 or IPv6, not both.
- The NVE source interface loopback for VTEP can be IPv4 (VXLANv4) or IPv6 (VXLANv6), but not both.
- Overlay next-hop addresses in BGP EVPN must be resolved in the underlay URIB to the same address family. For example, IPv4 VTEP addresses require IPv4 BGP EVPN peering.
- Usage of IPv6 LLA requires the TCAM Region for **ing-sup** to be re-carved from the default value of 512 to 768. This step requires a copy run start and reload.  
Using IPv6 Link-Local Address (LLA) requires a TCAM region ( **ing-sup** ) re-carve from 512 to 768, plus a copy run start and reload.
- vPC fabric peering is not supported with IPv6 multicast underlay.

## Supported platforms

*Table 1: Platforms and NX-OS releases supporting IPv6 underlay with VXLAN*

Platforms	Releases
<ul style="list-style-type: none"> <li>• Cisco Nexus 9300- EX, FX2, FX, and GX EOR switches</li> <li>• N9K-X97160YC-EX, N9K-X9736C-FX with N9K-C9508-FM-E2 fabric module</li> <li>• N9K-X97160YC-EX with N9K-C9508-FM-E fabric module</li> </ul>	9.3(1)

## VTEP (Leaf/Border) Support with IPv6 Underlay

*Table 2: VTEP function (leaf and border) support with IPv6 underlay*

Platforms	Releases
<ul style="list-style-type: none"> <li>• Cisco Nexus 9332C</li> <li>• Cisco Nexus 9364C</li> <li>• Cisco Nexus 9300-EX</li> <li>• Cisco Nexus 9300-FX</li> <li>• Cisco Nexus 9300-FX2</li> <li>• Cisco Nexus 9300-FX3</li> <li>• Cisco Nexus 9300-FXP</li> </ul>	9.3(1)
Cisco Nexus 9300-GX	9.3(6)

## Supported features

VXLAN with IPv6 in the Underlay (VXLANv6) supports these features:

- Address management
  - ARP suppression in the overlay
  - DHCP
  - IGMP Snooping in the overlay
- Operational
  - VXLAN OAM
- Platform capabilities
  - Guestshell support

- Storm Control for host ports (Access Side)
- vPC and fabric features
  - Virtual Port Channel (vPC) with VIP and PIP support

### Unsupported features

VXLAN with IPv6 in the Underlay (VXLANv6) does not support these features:

- Forwarding and replication functions
  - Downstream VNI
  - VXLAN Flood and Learn
  - Static ingress replication (IR)
  - Tenant Routed Multicast (TRM)
- Monitoring and visibility
  - Bidirectional Forwarding Detection (BFD)
  - Overlay IGMP Snooping
  - NetFlow
  - Sampled Flow (sFLOW)
- Platform and hardware integrations
  - Cisco DCNM integration
  - Fabric Extender (FEX) attached to a VXLAN-enabled switch.
- Multi-homing and redundancy
  - EVPN Multi-homing with Ethernet Segment (ES)
  - vPC Fabric Peering
  - VXLAN Multi-Site
- VXLAN Access Features
  - 802.1x
  - Port security
  - PVLAN
  - PV translation (Switching/Routing)
  - Q-in-Q with QinVNI
  - QinVNI and Selective QinVNI
- Cross-domain and protocol handoff

- Cross Connect
- Multiprotocol Label Switching (MPLS) and Locator/ID Separation Protocol (LISP) handoff
- Multicast protocol support
  - Multicast underlay (PIM-BiDir, PIM-ASM, Snooping)
- CLI and network functions
  - **peer vtep** command
  - VNF multipath

### Overlay and NVE loopback support

## Configure the VTEP IP address

Establish a valid VTEP IP address to enable VXLAN endpoint connectivity.

Follow these steps to configure the VTEP IP address:

#### Before you begin

- Ensure you have a loopback interface configured on the switch.
- Verify that the loopback interface uses a unique, non-link-local IPv6 address in /128 format.
- Confirm the IPv6 address is advertised in the transport network using a dynamic routing protocol.

#### Procedure

**Step 1** Enter global configuration mode.

##### Example:

```
switch# configure terminal
switch(config)#
```

**Step 2** Configure the NVE interface.

##### Example:

```
switch(config)# interface nve1
switch(config-if-nve)#
```

**Step 3** Specify the source interface as the loopback interface.

##### Example:

```
switch(config-if-nve)# source interface loopback 1
```

The source interface must be a loopback interface that is configured on the switch with a valid /128 IP address. This /128 IP address must be known by the intermediate devices in the transport network and the remote VTEPs. This is accomplished by advertising it through a dynamic routing protocol in the transport network.

**Note**

The IPv6 address on `loopback1` must be a /128 address.

The VTEP IP address cannot be a link local IPv6 address.

**Step 4** Exit configuration mode.

**Example:**

```
switch(config-if-nve) # exit
switch(config) #
```

**Step 5** Configure the loopback interface.

**Example:**

```
switch(config) # interface loopback 1
switch(config-if) #
```

**Step 6** Configure IPv6 address on the interface.

**Example:**

```
switch(config-if) # ipv6 address 2001:db8:0:0:1:0:0:1/128
```

**Step 7** Exit configuration mode.

**Example:**

```
switch(config-if) # exit
switch(config) #
```

---

The switch is now configured with a valid VTEP IP address for VXLAN operation.

## Configure vPC for VXLANv6

Configure a vPC for VXLAN with IPv6 in the underlay, enabling VTEP functionality with distinct VIP (Virtual IP) and PIP (Primary IP) addresses.

•

IPv6 underlay for VXLAN does not support secondary IP addresses as IPv4 does. Each loopback interface must be assigned a unique /128 IPv6 address. The VIP and PIP each require dedicated loopbacks, and correct use of the `anycast loopback` command ensures proper traffic routing and redundancy in a vPC setup.

Follow these steps to configure vPC for VXLAN with IPv6 in the underlay:

**Before you begin**

- Ensure that loopback interfaces used for VIP and PIP have globally routable /128 IPv6 addresses, not link-local addresses, and that these addresses are advertised in the transport network using a dynamic routing protocol.
- The `anycast loopback` command is used only with VXLAN over IPv6 in the underlay.

## Procedure

**Step 1** Enter global configuration mode.

**Example:**

```
switch# configure terminal
switch(config)#
```

**Step 2** Configure the NVE interface.

**Example:**

```
switch(config)# interface nve1
switch(config-if-nve)#
```

**Step 3** Specify the PIP source interface and VIP anycast loopback.

**Example:**

```
switch(config-if-nve)# source interface loopback 1 anycast loopback 2
```

The source interface must be a loopback interface that is configured on the switch with a valid /128 IP address. This /128 IP address must be known by the transient devices in the transport network and the remote VTEPs. This is accomplished by advertising it through a dynamic routing protocol in the transport network.

**Note**

The IPv6 address on loopback1, the primary IP address (PIP), and loopback2, the secondary IP address (VIP), must be a /128 address.

The VTEP IP address cannot be a link local IPv6 address.

**Step 4** Exit configuration mode.

**Example:**

```
switch(config-if-nve)# exit
switch(config)#
```

**Step 5** Configure the loopback interface for the PIP.

**Example:**

```
switch(config)# interface loopback 1
switch(config-if)#
```

a) Configure IPv6 address on the interface.

**Example:**

```
switch(config-if)# ipv6 address 2001:db8:0:0:1:0:0:1/128
```

**Step 6** Exit loopback interface configuration.

**Example:**

```
switch(config-if)# exit
switch(config)#
```

**Step 7** Configure the loopback interface for the VIP (anycast).

**Example:**

```
switch(config)# interface loopback 2
```

a) Configure IPv6 address on the interface.

**Example:**

```
switch(config-if)# ipv6 address 2001:db8:0:0:1:0:0:2/128
```

**Step 8** Exit loopback interface configuration.

**Example:**

```
switch(config-if)# exit
switch(config)#
```

The vPC is configured for VXLAN with IPv6 in the underlay, with separate /128 IPv6 addresses assigned for the PIP and VIP on individual loopbacks. The configuration enables high-availability VTEP operation across the vPC pair.

## VXLANv6 configuration requirements and examples

VXLANv6 configuration requires careful attention to BGP setup, router ID assignment, and support for L2VPN EVPN address-family sessions. The following reference provides direct configuration examples and key requirements.

### BGP IPv6 next-hop

**BGP IPv6 next-hop requirements:** With IPv6 addresses set or matched in the next-hop, BGP must set or match the IPv6 next-hop address for route type-2 (MAC-IP) and route type-5 (IP Prefix).

**Example route-map configuration:**

```
set ipv6 next-hop <vtep address>
match ipv6 next-hop <vtep address>
```

### BGP IPv6 neighbor and router ID

BGP IPv6 neighbor and router ID requirements:

- The BGP IPv6 neighbor must support the L2VPN EVPN address family.
- The router ID in VXLANv6 must be an IPv4 address (32-bit).
  - By default, NX-OS uses a loopback IPv4 address for the router ID. If no IPv4 address is present on a loopback, manually specify the router ID.
- The BGP Route Distinguisher (RD), which is 64 bits, can be configured using the AS number and an IPv4 address. For VXLANv6, configure the RD using an IPv4 address, similar to VXLAN IPv4 deployments.

**Sample BGP configuration for VXLANv6:**

```
feature bgp
nv overlay evpn

router bgp 64496
! IPv4 router id
```

```

router-id 192.0.2.0
! Redistribute the igp/bgp routes
address-family ipv6 unicast
    redistribute direct route-map allow

! For IPv6 session, directly connected peer interface
neighbor 2001:DB8:0:1::55
    remote-as 64496
    address-family ipv6 unicast

```

**Sample OSPFv3 configuration:**

```

feature ospfv3

router ospfv3 201
router-id 192.0.2.1

interface ethernet 1/2
ipv6 address 2001:0DB8::1/48
ipv6 ospfv3 201 area 0.0.0.10

```

**Sample IS-IS for IPv6 configuration:**

```

router isis Enterprise
is-type level-1
net 49.0001.0000.0000.0003.00

interface ethernet 2/1
ipv6 address 2001:0DB8::1/48
isis circuit-type level-1
ipv6 router isis Enterprise

```

## Verification commands for VXLANv6

This topic lists commands and their sample outputs to help you verify the status and configuration details for VXLAN with IPv6 in the underlay.

*Table 3: VXLANv6 verification commands*

Command	Purpose
<b>show running-config interface nve 1</b>	Displays interface NVE 1 running configuration information.
<b>show nve interface 1 detail</b>	Displays NVE interface detail.
<b>show nve peers</b>	Displays the peering time and VNI information for VTEP peers.
<b>show nve vni ingress-replication</b>	Displays NVE VNI ingress replication information.
<b>show nve peers 2018:1015::abcd:1234:3 int nv1 counters</b>	Displays NVE peers counter information.
<b>show bgp l2vpn evpn 1012.0383.9600</b>	Displays BGP L2VPN information for route type 2.

Command	Purpose
<b>show bgp l2vpn evpn 303:304::1</b>	Displays BGP L2VPN EVPN for route type 3.
<b>show bgp l2vpn evpn 5.116.204.0</b>	Displays BGP L2VPN EVPN for route type 5.
<b>show l2route peerid</b>	Displays L2route peerid.
<b>show l2route topology detail</b>	Displays L2route topology detail.
<b>show l2route evpn imet all detail</b>	Displays L2route EVPN imet detail.
<b>show l2route fl all</b>	Display L2route flood list detail.
<b>show l2route mac all detail</b>	Displays L2route MAC detail.
<b>show l2route mac-ip all detail</b>	Displays MAC address and host IP address.
<b>show ip route 1.191.1.0 vrf vxlan-10101</b>	Displays route table for VRF.
<b>show forwarding ipv4 route 1.191.1.0 detail vrf vxlan-10101</b>	Displays forwarding information.
<b>show ipv6 route vrf vxlan-10101</b>	Displays IPv6 routing table.
<b>show bgp l2vpn evpn</b>	Displays BGP's updated routes.
<b>show bgp evi <i>evi-id</i></b>	Displays BGP EVI information.
<b>show forwarding distribution peer-id</b>	Displays forwarding information.
<b>show forwarding nve l2 ingress-replication-peers</b>	Displays forwarding information for ingress replication.
<b>show forwarding nve l3 peers</b>	Displays nv3 Layer 3 peers information.
<b>show forwarding ecmp platform</b>	Displays forwarding ECMP platform information.
<b>show forwarding ecmp platform</b>	Displays forwarding ECMP platform information.
<b>show forwarding nve l3 ecmp</b>	Displays forwarding NVE Layer 3 ECMP information.

### Verification commands

Below are commonly used verification commands:

#### General interface and NVE status

- **show running-config interface nve 1**

```
switch# show running-config interface nve 1
interface nve1
  no shutdown
  source-interface loopback1 anycast loopback2
  host-reachability protocol bgp
  member vni 10011
```

```

    ingress-replication protocol bgp
    member vni 20011 associate-vrf

```

- **show nve interface 1 detail**

```

switch# show nve interface nve 1 detail
Interface: nve1, State: Up, encapsulation: VXLAN
VPC Capability: VPC-VIP-Only [notified]
Local Router MAC: a093.51cf.78f7
Host Learning Mode: Control-Plane
Source-Interface: loopback1 (primary: 30:3:1::2)
Anycast-Interface: loopback2 (secondary: 303:304::1)
Source Interface State: Up
Anycast Interface State: Up
Virtual RMAC Advertisement: Yes
NVE Flags:
Interface Handle: 0x49000001
Source Interface hold-down-time: 745
Source Interface hold-up-time: 30
Remaining hold-down time: 0 seconds
Virtual Router MAC: 0600.0000.0001
Interface state: nve-intf-add-complete

```

- **show nve peers**

```

switch# show nve peers
Interface Peer-IP           State LearnType Uptime   Router-Mac
-----
nve1      1:1::1:1                Up      CP        00:44:09 5087.89d4.6bb7

```

- **show nve vni ingress-replication**

```

switch# show nve vni ingress-replication
Interface VNI      Replication List Source Up Time
-----
nve1      10011      1:1::1:1          BGP-IMET  00:46:55

```

- **show nve peers ipv6-address int nv1 counters**

```

switch# show nve peers 2018:2015::abcd:1234:3 int nve 1 counters
Peer IP: 2018:1015::abcd:1234:3
TX
    0 unicast packets 0 unicast bytes
    0 multicast packets 0 multicast bytes
RX
    0 unicast packets 0 unicast bytes
    0 multicast packets 0 multicast bytes

```

- **show nve peers ipv6-address vni 20001 vni int nv1 counters**

```

switch# show nve peers
Peer IP: 2018:1015::abcd:1234:3
VNI: 20001
RX
    0 unicast packets 0 unicast bytes
    0 multicast packets 0 multicast bytes

```

## BGP L2VPN/EVPN verification

- **show bgp l2vpn evpn**

- For (Route-Type 2)

```

switch# show bgp l2vpn evpn 1012.0383.9600
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 30.3.1.1:34067 (L2VNI 2001300)
BGP routing table entry for [2]:[0]:[0]:[48]:[1012.0383.9600]:[0]:[0.0.0.0]/216,
version 1051240
Paths: (1 available, best #1)
Flags: (0x000102) (high32 00000000) on xmit-list, is not in l2rib/evpn
Multipath: iBGP

  Advertised path-id 1
  Path type: local, path is valid, is best path, no labeled nexthop
  AS-Path: NONE, path locally originated
    303:304::1 (metric 0) from 0:: (30.3.1.1)
      Origin IGP, MED not set, localpref 100, weight 32768
      Received label 2001300
      Extcommunity: RT:2:2001300 ENCAP:8

  Path-id 1 advertised to peers:
    2::21          2::66
BGP routing table entry for [2]:[0]:[0]:[48]:[1012.0383.9600]:[32]:[4.231.115.2]/272,
version 1053100
Paths: (1 available, best #1)
Flags: (0x000102) (high32 00000000) on xmit-list, is not in l2rib/evpn
Multipath: iBGP

  Advertised path-id 1
  Path type: local, path is valid, is best path, no labeled nexthop
  AS-Path: NONE, path locally originated
    303:304::1 (metric 0) from 0:: (30.3.1.1)
      Origin IGP, MED not set, localpref 100, weight 32768
      Received label 2001300 3003901
      Extcommunity: RT:2:2001300 RT:2:3003901 ENCAP:8 Router MAC:0600.0000.0001

  Path-id 1 advertised to peers:
    2::21          2::66

```

- For Route-Type 3

```

switch# show bgp l2vpn evpn 303:304::1
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 30.3.1.1:32769 (L2VNI 2000002)
BGP routing table entry for [3]:[0]:[128]:[303:304::1]/184, version 1045060
Paths: (1 available, best #1)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn
Multipath: iBGP

  Advertised path-id 1
  Path type: local, path is valid, is best path, no labeled nexthop
  AS-Path: NONE, path locally originated
    303:304::1 (metric 0) from 0:: (30.3.1.1)
      Origin IGP, MED not set, localpref 100, weight 32768
      Extcommunity: RT:2:2000002 ENCAP:8
      PMSI Tunnel Attribute:
        flags: 0x00, Tunnel type: Ingress Replication
        Label: 2000002, Tunnel Id: 303:304::1

  Path-id 1 advertised to peers:
    2::21          2::66

```

- For Route-Type 5

```

switch# show bgp l2vpn evpn 5.116.204.0
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 2.0.0.52:302
BGP routing table entry for [5]:[0]:[0]:[24]:[5.116.204.0]/224, version 119983

```

```

Paths: (2 available, best #2)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn, is not in
HW
Multipath: iBGP

Path type: internal, path is valid, not best reason: Neighbor Address, no labeled
nexthop
Gateway IP: 0.0.0.0
AS-Path: 65001 5300 , path sourced external to AS
3::52 (metric 200) from 2::66 (2.0.0.66)
Origin IGP, MED not set, localpref 100, weight 0
Received label 3003301
Extcommunity: RT:2:3003301 ENCAP:8 Router MAC:f80b.cb53.4897
Originator: 2.0.0.52 Cluster list: 2.0.0.66

Advertised path-id 1
Path type: internal, path is valid, is best path, no labeled nexthop
Imported to 2 destination(s)
Imported paths list: evpn-tenant-0301 default
Gateway IP: 0.0.0.0
AS-Path: 65001 5300 , path sourced external to AS
3::52 (metric 200) from 2::21 (2.0.0.21)
Origin IGP, MED not set, localpref 100, weight 0
Received label 3003301
Extcommunity: RT:2:3003301 ENCAP:8 Router MAC:f80b.cb53.4897
Originator: 2.0.0.52 Cluster list: 2.0.0.21

Path-id 1 not advertised to any peer

```

## L2 routes verification

### • show l2route peerid

```

switch# show l2route peerid
NVE Ifhdl IP Address PeerID Ifindex Num of
MAC's Num of NH's
-----
-----
1224736769 4999:1::1:1:1 4 1191182340 23377
0

```

### • show l2route topology detail

```

switch# show l2route topology detail
Flags: (L2cp)=L2 Ctrl Plane; (Dp)=Data Plane; (Imet)=Data Plane BGP IMET; (L3cp)=L3 Ctrl
Plane; (Bfd)=BFD over Vxlan; (Bgp)=BGP EVPN; (Of)=Open Flow mode; (Mix)=Open Flow IR
mixed mode; (Acst)=Anycast GW on spine;
Topology ID Topology Name Attributes
-----
-----
101 Vxlan-10101 VNI: 10101
Encap:1 IOD:0 IfHdl:1224736769
VTEP IP: 5001:1::1:1:7
Emulated IP: ::
Emulated RO IP: 0.0.0.0
TX-ID: 2004 (Rcvd Ack: 0)
RMAC: 00fe.c83e.84a7, VRFID: 3
VMAC: 00fe.c83e.84a7
VMAC RO: 0000.0000.0000
Flags: L3cp, Sub_Flags: --, Prev_Flags: -

```

### • show l2route evpn imet all detail

```
switch# show l2route evpn imet all detail
Flags- (F): Originated From Fabric, (W): Originated from WAN

Topology ID VNI Prod IP Addr Eth Tag PMSI-Flags Flags Type Label(VNI)
Tunnel ID NFN Bitmap
-----
901 10901 BGP 4999:1::1:1:1 0 0 - 6 10901
4999:1::1:1:1
```

- **show l2route fl all**

```
switch# show l2route fl all
Topology ID Peer-id Flood List Service Node
-----
901 4 4999:1::1:1:1 no
```

- **show l2route mac all detail**

```
switch# show l2route mac all detail

Flags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link
(Dup):Duplicate (Spl):Split (Rcv):Rcv (AD):Auto-Delete (D):Del Pending
(S):Stale (C):Clear, (Ps):Peer Sync (O):Re-Originated (Nho):NH-Override
(Pf):Permanently-Frozen, (Orp): Orphan

Topology Mac Address Prod Flags Seq No Next-Hops
-----
901 0016.0901.0001 BGP SplRcv 0 6002:1::1:1:1

Route Resolution Type: Regular
Forwarding State: Resolved (PeerID: 2)
Sent To: L2FM
Encap: 1
```

- **show l2route mac-ip all detail**

```
switch# show l2route mac-ip all detail
Flags -(Rmac):Router MAC (Stt):Static (L):Local (R):Remote (V):vPC link
(Dup):Duplicate (Spl):Split (Rcv):Rcv (D):Del Pending (S):Stale (C):Clear
(Ps):Peer Sync (Ro):Re-Originated (Orp):Orphan

Topology Mac Address Host IP Prod Flags
Seq No Next-Hops
-----
901 0016.0901.0001 46.1.1.101 BGP --
0 6002:1::1:1:1
Sent To: ARP
encap-type:1
```

## IP and IPv6 route verification

- **show ip route 1.191.1.0 vrf vxlan-10101**

```
switch# show ip route 1.191.1.0 vrf vxlan-10101
IP Route Table for VRF "vxlan-10101"
'*' denotes best ucast next-hop
'***' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

1.191.1.0/29, ubest/mbest: 6/0
```

```

    *via fe80::2fe:c8ff:fe09:8fff%default, Po1001, [200/0], 00:56:21, bgp-4002, internal,
    tag 4007 (evpn)
    segid: 10101 VTEP:(5001:1::1:1:1, underlay_vrf: 1) encap: VXLAN

    *via fe80::2fe:c8ff:fe09:8fff%default, Po1002, [200/0], 00:56:21, bgp-4002, internal,
    tag 4007 (evpn)
    segid: 10101 VTEP:(5001:1::1:1:1, underlay_vrf: 1) encap: VXLAN

    *via fe80::2fe:c8ff:fe09:8fff%default, Po1001, [200/0], 00:56:32, bgp-4002, internal,
    tag 4007 (evpn)
    segid: 10101 VTEP:(5001:1::1:1:2, underlay_vrf: 1) encap: VXLAN

    *via fe80::2fe:c8ff:fe09:8fff%default, Po1002, [200/0], 00:56:32, bgp-4002, internal,
    tag 4007 (evpn)
    segid: 10101 VTEP:(5001:1::1:1:2, underlay_vrf: 1) encap: VXLAN

```

- **show forwarding ipv4 route 1.191.1.0 detail vrf vxlan-10101**

```

switch# show forwarding ipv4 route 1.191.1.0 detail vrf vxlan-10101

slot 1
=====
Prefix 1.191.1.0/29, No of paths: 2, Update time: Mon Apr 15 15:38:17 2019

    5001:1::1:1:1      nve1
    5001:1::1:1:2      nve1

```

- **show ipv6 route vrf vxlan-10101**

```

switch# show ipv6 route vrf vxlan-10101
IPv6 Routing Table for VRF "vxlan-10101"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]

2:2:2::101/128, ubest/mbest: 1/0
    *via 5001:1::1:1:1/128%default, [200/0], 00:55:31, bgp-4002, internal, tag 4002
(evpn) segid 10101
VTEP:(5001:1::1:1:1, underlay_vrf: 1) encap: VXLAN

```

## Forwarding verification

- **show forwarding distribution peer-id**

```

switch# show forwarding distribution peer-id
UFDM Peer-id allocations: App id 0
App: VXLAN   Vlan: 1      Id: 4999:1::1:1:1 0x49030001 Peer-id: 0x6
App: VXLAN   Vlan: 1      Id: 5001:1::1:1:1 0x49030001 Peer-id: 0x2
App: VXLAN   Vlan: 1      Id: 5001:1::1:1:2 0x49030001 Peer-id: 0x1
App: VXLAN   Vlan: 1      Id: 5001:1::1:1:7 0x49030001 Peer-id: 0x7
App: VXLAN   Vlan: 1      Id: 5001:1::1:2:101 0x49030001 Peer-id: 0x8
App: VXLAN   Vlan: 1      Id: 5001:1::1:2:102 0x49030001 Peer-id: 0x5
App: VXLAN   Vlan: 1      Id: 5001:1::1:2:103 0x49030001 Peer-id: 0x9
App: VXLAN   Vlan: 1      Id: 5001:1::1:2:104 0x49030001 Peer-id: 0xa
App: VXLAN   Vlan: 1      Id: 5001:1::1:2:105 0x49030001 Peer-id: 0xb
App: VXLAN   Vlan: 1      Id: 5001:1::1:2:106 0x49030001 Peer-id: 0xc
App: VXLAN   Vlan: 1      Id: 5001:1::1:2:107 0x49030001 Peer-id: 0xd

```

- **show forwarding nve l2 ingress-replication-peers**

```

switch# show forwarding nve l2 ingress-replication-peers
slot 1
=====

```

```

Total count of VLANs with ingr-repl peers: 1950
VLAN 1024 VNI 0 Vtep Ifindex 0x0 plt_space : 0x1ca75e14
  peer : 6002:1::1:1:1
  peer : 5001:1::1:1:7
  peer : 4999:1::1:1:1

PSS VLAN:1024, VNI:0, vtep:0x0x0, peer_cnt:3
  peer : 6002:1::1:1:1 marked : 0
  peer : 5001:1::1:1:7 marked : 0
  peer : 4999:1::1:1:1 marked : 0
VLAN 1280 VNI 0 Vtep Ifindex 0x0 plt_space : 0x1ca75e14
  peer : 6002:1::1:1:1
  peer : 5001:1::1:1:7
  peer : 4999:1::1:1:1

PSS VLAN:1280, VNI:0, vtep:0x0x0, peer_cnt:3
  peer : 6002:1::1:1:1 marked : 0
  peer : 5001:1::1:1:7 marked : 0
  peer : 4999:1::1:1:1 marked : 0

```

### • show forwarding nve l3 peers

```

switch# show forwarding nve l3 peers
slot 1
=====

```

```

EVPN configuration state: disabled, PeerVni Adj enabled
NVE cleanup transaction-id 0

```

tunnel_id	Peer_id	Peer_address	Interface	rmac	origin	state	del	count
0x0	1225261062	4999:1::1:1:1	nve1	0600.0001.0001	URIB	merge-done		
no	100							
0x0	1225261058	5001:1::1:1:1	nve1	2cd0.2d51.9f1b	NVE	merge-done		
no	100							
0x0	1225261057	5001:1::1:1:2	nve1	00a6.cab6.bbbb	NVE	merge-done		
no	100							
0x0	1225261063	5001:1::1:1:7	nve1	00fe.c83e.84a7	URIB	merge-done		
no	100							
0x0	1225261064	5001:1::1:2:101	nve1	0000.5500.0001	URIB	merge-done		
no	100							
0x0	1225261061	5001:1::1:2:102	nve1	0000.5500.0002	URIB	merge-done		
no	100							
0x0	1225261065	5001:1::1:2:103	nve1	0000.5500.0003	URIB	merge-done		
no	100							
0x0	1225261066	5001:1::1:2:104	nve1	0000.5500.0004	URIB	merge-done		
no	100							
0x0	1225261067	5001:1::1:2:105	nve1	0000.5500.0005	URIB	merge-done		
no	100							

### ECMP verification

#### • show forwarding ecmp platform

```

switch# show forwarding ecmp platform
slot 1
=====
ECMP Hash: 0x198b8aae, Num Paths: 2, Hw index: 0x17532
Partial Install: No
Hw ecmp-index: unit-0:1073741827 unit-1:0 unit-2:0, cmn-index: 95538
Hw NVE ecmp-index: unit-0:0 unit-1:0 unit-2:0, cmn-index: 95538
Refcount: 134, Holder: 0x0, Intf: Ethernet1/101, Nex-Hop: fe80:7::1:2
  Hw adj: unit-0:851977 unit-1:0 unit-2:0, cmn-index: 500010 LIF:4211
  Intf: Ethernet1/108, Nex-Hop: fe80:8::1:2
  Hw adj: unit-0:851978 unit-1:0 unit-2:0, cmn-index: 500012 LIF:4218

```

```
VOBJ count: 0, VxLAN VOBJ count: 0, VxLAN: 0
```

```
ECMP Hash: 0x2bb2905e, Num Paths: 3, Hw index: 0x17533
Partial Install: No
Hw ecmp-index: unit-0:1073741828 unit-1:0 unit-2:0, cmn-index: 95539
Hw NVE ecmp-index: unit-0:0 unit-1:0 unit-2:0, cmn-index: 95539
Refcunt: 16, Holder: 0x0, Intf: Ethernet1/101, Nex-Hop: fe80:7::1:2
  Hw adj: unit-0:851977 unit-1:0 unit-2:0, cmn-index: 500010 LIF:4211
  Intf: Ethernet1/108, Nex-Hop: fe80:8::1:2
  Hw adj: unit-0:851978 unit-1:0 unit-2:0, cmn-index: 500012 LIF:4218
  Intf: port-channel1003, Nex-Hop: fe80:9::1:2
  Hw adj: unit-0:851976 unit-1:0 unit-2:0, cmn-index: 500011 LIF:4106
VOBJ count: 0, VxLAN VOBJ count: 0, VxLAN: 0
```

- switch# **show forwarding ecmp recursive**

```
slot 1
=====
```

```
Virtual Object 17 (vxlan):
  Hw vobj-index (0): unit-0:851976 unit-1:0 unit-2:0, cmn-index: 99016
  Hw NVE vobj-index (0): unit-0:0 unit-1:0 unit-2:0, cmn-index: 99016
  Hw vobj-index (1): unit-0:0 unit-1:0 unit-2:0, cmn-index: 0
  Hw NVE vobj-index (1): unit-0:0 unit-1:0 unit-2:0 cmn-index: 0
  Num prefixes : 1
Partial Install: No
Active paths:
  Recursive NH 5001:1::1:2:10a/128 , table 0x80000001
CNHs:
  fe80:9::1:2, port-channel1003
  Hw adj: unit-0:851976 unit-1:0 unit-2:0, cmn-index: 500011, LIF:4106
  Hw NVE adj: unit-0:0 unit-1:0 unit-2:0, cmn-index: 500011, LIF:4106
  Hw instance new : (0x182c8, 99016) ls count new 1
  FEC: fec_type 0
  VOBJ Refcount : 1
Virtual Object 167 (vxlan): ECMP-idx1:0x17536(95542), ECMP-idx2:0x0(0),
  Hw vobj-index (0): unit-0:1073741832 unit-1:0 unit-2:0, cmn-index: 99166
  Hw NVE vobj-index (0): unit-0:3 unit-1:0 unit-2:0, cmn-index: 99166
  Hw vobj-index (1): unit-0:0 unit-1:0 unit-2:0, cmn-index: 0
  Hw NVE vobj-index (1): unit-0:0 unit-1:0 unit-2:0 cmn-index: 0
  Num prefixes : 1
Partial Install: No
Active paths:
  Recursive NH 5001:1::1:3:125/128 , table 0x80000001
CNHs:
  fe80:7::1:2, Ethernet1/101
  Hw adj: unit-0:851977 unit-1:0 unit-2:0, cmn-index: 500010, LIF:4211
  Hw NVE adj: unit-0:0 unit-1:0 unit-2:0, cmn-index: 500010, LIF:4211
  fe80:8::1:2, Ethernet1/108
  Hw adj: unit-0:851978 unit-1:0 unit-2:0, cmn-index: 500012, LIF:4218
  Hw NVE adj: unit-0:0 unit-1:0 unit-2:0, cmn-index: 500012, LIF:4218
  Hw instance new : (0x1835e, 99166) ls count new 2
  FEC: fec_type 0
  VOBJ Refcount : 1
```

- **show forwarding nve l3 ecmp**

```
switch# show forwarding nve l3 ecmp
slot 1
=====
  ECMP Hash: 0x70a50e4, Num Paths: 2, Hw Index: 0x17534
  table_id: 403, flags: 0x0, adj_flags: 0x0, Ref-ct: 101
  tunnel_id: 5001:1::1:1:1, segment_id: 10101
  tunnel_id: 5001:1::1:1:2, segment_id: 10101
  Hw ecmp-index: unit0: 1073741830 unit1: 0 unit2: 0
```

```
ECMP Hash: 0x1189f35e, Num Paths: 2, Hw Index: 0x17535
table_id: -2147483245, flags: 0x0, adj_flags: 0x0, Ref-ct: 50
  tunnel_id: 5001:1::1:1:1, segment_id: 10101
  tunnel_id: 5001:1::1:1:2, segment_id: 10101
Hw ecmp-index: unit0: 1073741831 unit1: 0 unit2: 0
```

